Are we moving to a less toxic world? How bans on Persistent Organic Pollutants affect firms' inventive strategies

Gianluca Biggi* Elisa Giuliani† Arianna Martinelli‡ Julia Mazzei§

PRELIMINARY DRAFT: DO NOT CITE OR QUOTE

Abstract

Over the years, our economic and innovation systems have harnessed chemistry to exploit its potential for private welfare and profit, resulting in groundbreaking advancements in synthetic chemistry after World War II. However, the diffusion of these technologies has revealed significant unintended consequences, leading to detrimental impacts on human health and ecosystems. The hazards of a class of pesticides, known as persistent organic pollutants (POPs), were brought to light as early as 1962. Starting the same year, several countries have legislated with the aim of phasing out the production, use, import, and export of POPs. This regulatory push was expected to reduce firms' innovative efforts in this area. However, path dependency and existing investments in production facilities may have hindered this process, leading to cumulative innovation and attempts to invent around the bans. This paper seeks to investigate the impact of POPs regulation on firms' innovation trajectories, determining to what extent the ban redirect the research efforts toward less toxic domains. By combining several unique datasets, we apply a firm-level Difference-in-differences with multiple time periods analysis, leveraging the staggered timing of the bans. Our findings reveal that country-level initiatives aimed at banning POPs were effective in reducing patenting activities in the domain of POPs only in chemical fields, while the number of POPs-related patents filed outside the chemical fields has increased despite the ban.

Keywords: chemical inventions, persistent organic pollutants (POPs), policy evaluation, patent toxicity

JEL classification: Q55, Q58, O31, O33

*Institute of Economics and EMbeDS, Scuola Superiore Sant'Anna, Pisa (Italy). Corresponding author: gianluca.biggi@santannapisa.it

†Department of Economics and Management, University of Pisa (Italy).

‡Institute of Economics and EMbeDS, Scuola Superiore Sant'Anna, Pisa (Italy).

§Institute of Economics and EMbeDS, Scuola Superiore Sant'Anna, Pisa (Italy).

1 Introduction

Planetary boundaries represent a framework for understanding and quantifying the environmental limits within which humanity can safely operate to maintain a stable and sustainable planet (Rockström et al., 2009). One core pillar of the planetary boundaries is chemical pollution, which considers the release of synthetic chemical pollutants into the environment. A key concern is represented by the introduction of harmful chemicals that can have widespread and long-lasting effects on ecosystems and human health, also called persistent organic pollutants (POPs). POPs are "chemical substances that persist in the environment, bio-accumulate through the food chain, and pose a risk of causing adverse effects to human health and the environment" (UNEP, 2017, p. 6). The reason why there is so much concern is that they are highly hazardous: they cause cancer, damages to the immune and nervous systems, and are highly poisonous to human beings, animal species and the environment at large (AMAP, 2009). Evidence of the dangers of a class of a first group of hazardous POPs, also known as the "Dirty Dozen", started to be available already in 1962 (Carson, 2002)[1], yet it took some forty years for governments to agree on a global ban. DDT is probably the most infamous POP; while it has contributed to eradicating malaria and promoting economic development (Özkara et al., 2016), its noxious effects have been known quite early.

POPs have been globally banned under the Stockholm Convention, promoted by the United Nations, which has led to the ban of a first group of 12 compounds in 2001 and the second ban of an additional 16 compounds in 2017.[2] Country-level initiatives have anticipated this global effort started in 1960 and aimed at phasing out the production of POPs. While not homogeneous across countries, such country-level regulations have banned either the use, import or export of all or some POPs. As these bans aim to regulate POPs, we expect that firms active in this area would lower their innovative effort around these toxic compounds. However, path dependency and investments in production facilities might hamper this process and trigger a process of cumulative innovation and "inventing around" to circumvent the ban.

This paper aims to identify the impact of the regulation of POPs on firms' innovation trajectory and assess to what extent the ban has moved firms' research efforts towards new domains that are less toxic. We rely on two unique and novel datasets to tackle this research agenda. The first dataset reconstructs the inventive activities of chemical firms. We exploit the chemical database CAS SciFinderN to retrieve all the patents containing at least one POPs chemical compound, identified through their international chemical identifiers (i.e., the CAS Registry Number). We then use the EPO-PATSTAT database and

[1] For more recent evidence, see: Bartrons et al. (2016), Jepson and Law (2016), Ma et al. (2011).

[2] The first 12 banned POPs, also known as "Dirty Dozen", are aldrin, chlordane, dieldrin, endrin, heptachlor, hexachlorobenzene (HCB), mirex, toxaphene, polychlorinated biphenyls (PCBs), dichlorodiphenyltrichloroethane (DDT), dioxin, and polychlorinated dibenzofurans. The subsequent 16 banned compounds are chlordecone, α -hexachlorocyclohexane (α -HCH) and β-hexachlorocyclohexane (β-HCH), hexabromodiphenyl ether (hexaBDE) and heptabromodiphenyl ether (heptaBDE), lindane, pentachlorobenzene (PeCB), tetrabromodiphenyl ether (tetra BDE) and pentabromodiphenyl ether (pentaBDE), perfluorooctanesulfonic acid (PFOS), endosulfans, and hexabromocyclododecane (HBCD)

ORBIS-IP to obtain standard patent-related information for each identified patent and applicant. The second hand-collected dataset includes information on country-specific regulations on POPs since 1960. This information includes the type of restriction imposed on each POP (e.g. ban on the use, import, export and production and related exemptions when applied) along with the year of implementation.

Combining these two datasets we apply a Difference-in-differences with multiple time periods analysis (Callaway and Sant'Anna, 2021), exploiting the staggered timing of the bans and considering that the same firm can protect inventions in different countries. This research design allows us to precisely identify the impact of banning POPs on firms' innovation activity and possibly draw causal relations. This methodology is particularly suited to our study because the timing of the bans is staggered. Our identification strategy relies on the firms' widespread practice to fill applications of worthy inventions in multiple patent offices worldwide to obtain patent protection in multiple countries (de Rassenfosse et al., 2021). This implies that we can observe the inventive strategies of a given firm in countries with different regulatory settings.

We use several outcome variables computed at the firm-level to proxy firms' inventive strategies. First, we rely on the number of POPs patent applications filed by the firm in each country-year. Second, we compute the similarity of the firm's inventions over time based on the indicator of chemical compounds' novelty proposed by Krieger et al. (2022). Third, in line with Biggi et al. (2022), we measure the degree of toxicity of patents considering all the chemical structures therein. For computational reasons, comprehensive regression results are reported only using the first dependent variable at this stage.

Our preliminary findings suggest that banning POPs has been effective in reducing patenting activities in the domain of POPs only in chemical fields. In contrast, POPs-related patents filed outside the chemical fields have increased over time despite the ban.

Moreover, to map the evolution of technological trajectories in the domain of POPs, chemical data allows us to apply a recently developed measure of the similarity of chemical inventions (Krieger et al., 2022), going beyond established indicators used in the innovation literature. Second, we provide a policy evaluation of the numerous countries' initiatives worldwide aimed at banning the use of POPs. To the best of our knowledge, no existing studies provide an empirical assessment of the effectiveness of these policies. Understanding the economic impact of chemical regulations and how they influence firms' innovation trajectories is important for public policy, R&D management and firm strategy.

2 Background

2.1 Countries' legislative initiatives that led to the Stockholm Convention

POPs are a class of pollutants characterized by their persistence in the environment and bioaccumulation through the food chain (UNEP, 2017). Such pollutants are highly hazardous: they cause cancer, damages

the immune and nervous systems, and are highly poisonous to human beings, animal species and the environment at large (AMAP, 2009). From the economic perspective, they are a global negative externality as they can be detected in places where they are not directly used. For example, a recently growing concern expressed by bioscientists concerns the re-volatilization POPs as a result of the melting of polar ice (Nizzetto et al., 2010; Rigét et al., 2019). This group of pollutants consists mainly of pesticides (e.g. DDT), industrial chemicals and unintentional by-products of industrial processes. Because of the increasing evidence of the danger of these chemicals and following the recommendation of the Intergovernmental Forum on Chemical Safety (IFCS) and the International Programme on Chemical Safety (IPCS), a first group of a hazardous POPs, also known as the "Dirty Dozen"(namely: Aldrin, Chlordane, DDT, Endrin, Heptachlor, Hexachlorobenzene, Mirex, Pentachlorobenzene, Toxaphene, and Polychlorinated Biphenyls; Polychlorinated dibenzo-p- dioxins), were banned or their use was severely restricted starting from 2001 by an international treaty – the Stockholm Convention. The treaty was signed between 2001 and 2002 by 153 countries worldwide, and it was progressively ratified in the following years (most ratifications occurred between 2001 and 2007). This global ban has mitigated concerns about these hazardous chemicals, at least up until very recently, when it became clear that they would re-volatilize in the atmosphere as a consequence of melting Antarctic ice (Nizzetto et al., 2010; Ma et al., 2011; Rigét et al., 2019). In 2017, the Stockholm Convention added 16 new POPs to the treaty and currently has other chemical compounds under review, including per-polyfluoroalkyl substances (PFAS). More recently, Cousins et al. (2022) shows that environmental contamination caused by PFAS is pushing chemical pollution outside the limits of planetary boundaries that define a safe and operating space for humanity (Rockström et al., 2009).

While it took some forty years for governments to agree on a global ban, evidence of the dangers of a class of pesticides started to be available in 1962 (Carson, 2002). In the same year, well before the global initiative of the Stockholm Convention, several countries have legislated, intending to phase out the production of POPs. For instance, the US has implemented several measures that mirror the treaty's regulations such as the Toxic Substances Control Act (TSCA), enacted in the 1970s and amended in 2016 that regulate the manufacture, distribution, use, and disposal of chemicals, including some POPs or the Clean Water Act (CWA), first enacted in 1972 that regulates discharges of toxic pollutants including POPs into US waters.

Pre-Stockholm Convention initiatives are rather heterogeneous across countries and differ across two main dimensions. First, while most countries regulate all POPs simultaneously, some bans focus on one or a group of POPs. Second, not all bans are equal in phasing out POPs, as they may affect these compounds' use, production, import or export differently. For instance, in the early 1970s, the US banned the agricultural use of DDT. The ban was primarily driven by environmental concerns about DDT's persistence and its impact on wildlife, particularly birds like the Bald Eagle and the Peregrine

Falcon, as well as the accumulation of DDT residues in the food chain. On the other hand, Australia through the National Pollution Inventory (NPI) of 1970s focused from the beginning on regulating all POPs.

However, while most research has addressed the legacy of POPs, very little is known about whether these country-level initiatives stemming in the '60 affected firms' innovation activity. For instance, to what extent do such regulatory initiatives affect a firm's innovative trajectory around POPs? Are firms responding to the ban moving towards less harmful innovation or just inventing around the ban? This paper aims to answer these research questions.

2.2 Firms' responses to banned technologies: Competing working hypotheses

While we expect products that have been banned to disappear from the market progressively, it is not clear what happens to the knowledge that underlies such products. Earlier research has shown that companies know more than they make (Brusoni et al., 2001), which means that companies tend to retain or produce knowledge in excess to cope with unpredictable changes in their competitive environment. In the context of toxic and banned technologies, we know virtually nothing about whether and why, in the face of a ban, companies and other organizations generate knowledge conducive to developing potentially hazardous new products. Unlike countries that invest in hazardous technologies – such as chemical or nuclear weapons – to deter more powerful adversaries from using them for destructive purposes, companies and other organizations are not likely moved by deterrence. However, knowledge about their patenting strategies in the aftermath of a ban is crucial to inform regulators about potential future hazards and deepen our understanding of corporate inventive and innovative processes in the context of sustainability transitions.

We develop two alternative general hypotheses about the patenting behaviour of companies and other organizations after the emergence of a ban. Both our hypotheses assume that companies and other organizations researching the banned technologies are not malevolent entities and will, therefore, have agreed on the threats posed by their past discoveries and will seek to find ways to address them in the future, to stop or minimize the noxious impacts on humans and the environment. The direction of their R&D efforts after the ban is also assumed to result from calculated strategic decisions in the face of a changing regulatory environment. Such decisions are expected to be taken to both address social and environmental threats and to hedge against future losses in the case of more stringent regulatory measures. In that context, our first hypothesis suggests that, after a ban, companies and other inventing entities will take new "explorative" innovation paths (March, 1991) and, therefore, will make the knowledge underpinning the banned technologies less important. They will potentially explore radically new technological trajectories (Dosi, 1982), which we envisage will not build on the established and "toxic" banned technological knowledge, which will instead decline in relevance. The general rationale for this is

that the banned technological knowledge becomes a dangerous terrain for new investments, especially if the new technological developments draw – even if partially – on the banned technological knowledge that can be subject to future bans within the same or a different treaty. Hence, in a scenario of heightened uncertainty and risk following a ban, we expect the following hypothesis to be supported:

Hypothesis 1 (a): After banning a given technology, the patents underpinning the banned technology will decline.

However, an alternative hypothesis is also possible. Changing technological trajectories and abandoning a beaten track can also be risky for inventing organizations, which may instead be willing to continue building on their existing knowledge basis rather than looking for radically new solutions. Path dependency and investments in production facilities might hamper this process and trigger a process of cumulative innovation and "inventing around" to circumvent the ban. Inventing organizations will continue to invest incrementally based on their existing knowledge base following an "exploitative" invention path (March, 1991). Unlike Hypothesis 1 (a), we would not observe a decline in patents connected to the banned technology. However, we expect new patents to be dissimilar and have a less toxic potential on human beings and the environment than the one they had before the ban. Hence, we formulate the following alternative hypothesis:

Hypothesis 1 (b): After banning a given technology, the patents underpinning the banned technology will not decline, while their similarity or degree of toxicity will.

3 Data and descriptive evidence

3.1 Building the chemical dataset

To test our hypothesis, we rely on various data sources (see Table A1 for an overview). First, we use the chemical database maintained by the Chemical Abstract Society (CAS) SciFinderN to retrieve all patent documents containing at least one chemical compound associated with a POP. The SciFinderN database has the unique feature of providing comprehensive patent-compounds associations. Therefore, we can precisely retrieve the POP-related patents using their compound international identifier (i.e. the CAS Registry Number). Furthermore, this dataset includes information about each compound's function in the invention (e.g., part of a mixture, catalyst, excipient, reagent). This information allows us to select only those patents building upon POPs - among other chemical compounds - as components of the products or processes sought of the invention and to exclude inventions aimed at removing or mitigating the harmful effect of POPs. The search conducted in SciFinderN resulted in 3,555 worldwide POPs patent families filed between 1960 and 2016.[3] We then match the patent identities retrieved from SciFinderN with the EPO-PATSTAT database to obtain standard patent-related information such as the filing date,

[3]Note that the SciFinderN provides all the relevant chemical data at the patent family level. However, each family is identified with only one family member, generally the first one applied.

applicant name and IPC classes. More importantly, for our identification strategy, we use the EPO-PATSTAT database to reconstruct the DOCDB family and identify all the legislations in which each POP-related invention is protected. We use the IPC classes and the classification by Van Looy et al. (2015) to distinguish patents filed in the chemical and non-chemical technological fields. Interestingly, the majority of POP patents are applied outside the chemical domain. Indeed, only 907 patents have the main IPC class that refers to chemical technological fields. This finding indicates that such harmful compounds are used for producing high-quality parts for manufacturing. For example, the metal plating process uses several additives in producing machinery and vehicles. One is a POP called perfuorooctane-sulfonic acid or –sulfonate (PFOS).[4]

3.2 Data on firms and countries' legislation

We devote particular effort to harmonizing and cleaning applicants' names to perform our firm-level analysis. Starting from the harmonized names provided by Magerman et al. (2006)[5], we manually cleaned more than 85% of the applicants' names. At this stage, we aggregate subsidiaries to single names to avoid duplicating the same entity. Our final dataset contains 898 distinct patent applicants, among which 695 are private firms, and the remaining 203 are research institutions, universities or public authorities (see Table A1). We complement applicants' information using ORBIS-IP, from which we retrieve some firm-level information.

Finally, for implementing the staggered Difference-in-Difference model, we compile a dataset including legislative information for each POP for 146 countries from 1960 onwards. Among these, we restrict our current analysis to 85 countries for which we have i) information about national patenting activity and ii) at least one POP patent filed in the period of analysis. Table 1 indicates the country included in the analysis and the year of the ban before the global ban of the Stockholm convention. The first country-level initiative goes back to 1962, when the Netherlands started legislating about POPs, followed by Sweden, Norway and Hungary in 1966, Estonia in 1967 and Argentina in 1968.

Figure 1 display the data reported in table 1 to give a better glimpse of the time heterogeneity of ban adoption. Such heterogeneity motivates the choice to implement a staggered Difference-in-Difference model.

3.3 Descriptive evidence

In Table 2, we display the top 20 applicants by the number of POPs patents filed throughout the period of analysis.

Bayer stands out as the company with the highest number of patents in the POPs domain, with 729 patents filed worldwide in the analysis period. In terms of patent shares in the overall period, Bayer,

[4]See https://www.unido.org.

[5]The EPO-PATSTAT database provides the standardized names in table TLS206_person.

Table 1: List of countries and year of ban

Country	Year of ban
NL	1962
NO, SE, HU	1966
EE	1967
AR	1968
BG, KR, DZ	1969
US, DK, FI, MD, TJ	1970
JP, IN, HR, TR, RS	1971
GB, DE	1972
BE, ID, GT	1974
LT	1975
CA, CY	1976
AU, CZ	1977
UY	1978
JO, TN	1980
TH, LU, IE	1981
ZA, CL, SI	1982
PH	1983
MA	1984
SK, MW, EC, PT, SG, MC, CO, ZM	1985
CH, ES, KE	1986
NZ	1987
AT, RO, MX, CR	1988
CU	1989
FR, HN, PE	1991
VN	1992
LV	1993
IS	1996
PA	1997
PL	2001
CN, BR, HK, DO, EG	2004
GR	2006
NI, ZW, UA	2007
RU, ME	2011
MT	2017
IT	2022
IL, MY, GE, SV, BA, TW, SM	Never

The year represents the first country's initiative in the domain of POPs. Datasource: hand-collected dataset.

Table 2: Top 20 applicants by number of POPs patents

Applicant	nr of POPs patents	POPs patent share	Country
Bayer	729	0.179	Germany
Basf	543	0.134	Germany
Syngenta	369	0.091	Switzerland
DuPont	153	0.038	United States
Dow Chemical	52	0.013	United States
FMC	49	0.012	United States
Zeneca	37	0.009	United Kingdom
MIT	33	0.008	United States
Ugine Kuhlmann	33	0.008	France
Imperial Chem	32	0.008	United Kingdom
Hoechst	30	0.007	Germany
Mitsubishi	26	0.006	Japan
Sanofi (Hungary)	22	0.005	Hungary
Aventis	21	0.005	France
Atochem	21	0.005	France
Spolchemie	21	0.005	Czech Republic
Albemarle	20	0.005	United States
Hooker Chemical	17	0.004	United States
Ishihara Sangyo Kaisha	17	0.004	Japan
Brigham&Women's (Harvard)	16	0.004	United States

Figure 1: Countries' initiatives against POPs over time

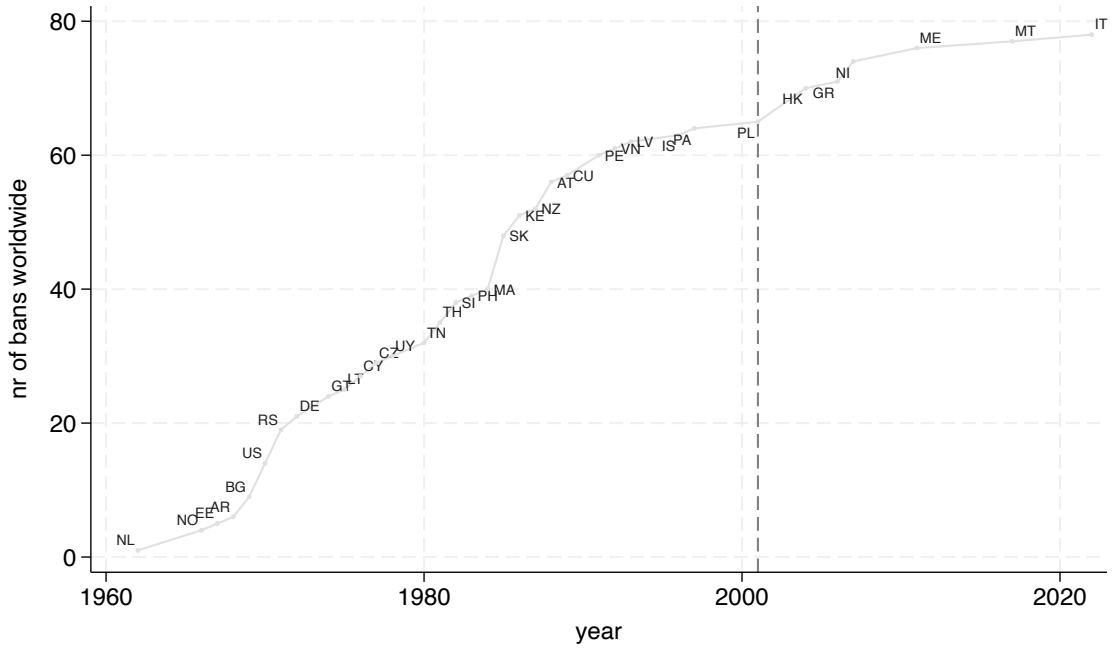

Our own elaboration on patents filed in 85 countries between 1962 and 2022. The y-axis is the cumulative number of countries in which either the production, import, export and use of POPs has been banned. The x-xis reports the year of the first country's initiative in the domain of POPs. The vertical line stands for the year in which the work of the Stockholm Convention starts (2001). Datasource: hand-collected dataset.

Basf and Syngenta hold 40% of the POPs patents. The list of top 20 applicants also includes universities and research centres such as the MIT.

In Figure 2, we plot the trend over time of the number of POPs patents filed worldwide between 1962 and 2016, together with the cumulative number of existing bans in the 85 countries analyzed. Surprisingly, despite the constant increase of countries' initiatives to ban POPs, the number of POP patents has remained stable until 2000 and has enormously increased since 2004. This trend is confirmed once looking at each country separately, as shown in Figure A1. While the surge in the number of POPs-related patents started in 2004 in most countries, a considerable increase in the number of applications of POPs-related inventions can be observed in the early 2000 in Japan.

4 Empirical strategy

4.1 Outcome variables: number of patents, similarity and toxicity

We use three different outcome variables to identify the impact of banning POPs on firms' innovation activity. First, we compute the number of patent applications, including a POP compound filed by the firm in each country and year. This variable captures the extensive margin of firms' innovation response to country-level legislation aimed at banning POPs.

Figure 2: Evolution of POPs-related patent applications and number of existing countries' regulation

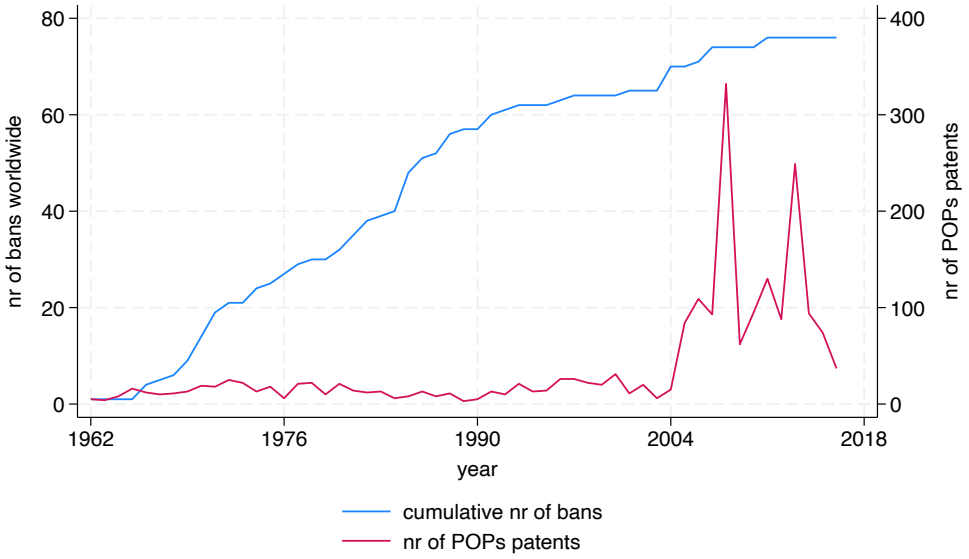

t

Our own elaboration on patents filed in 85 countries between 1962 and 2016. The left y-axis is the cumulative number of countries in which either the production, import, export and use of POPs has been banned. The right y-axis is the number of POPs patents filed worldwide in each year.

Second, we compute an ex-ante patent similarity measure by re-adapting the approach proposed by Krieger et al. (2022) to the chemical context.[6] This measure exploits the "Similarity Property Principle", which states that structurally similar compounds tend to have similar functional properties (Johnson et al., 1990). In this framework, the Tanimoto coefficient is a widely-used standard similarity measure in computational chemistry and drug discovery (Bajusz et al., 2015). The Tanimoto coefficient determines whether two or more compounds are closely related in terms of their structural characteristics, by considering a chemical compound as a set of structural features (i.e., the substructures or chemical fragments). Appendix B reports an example of how this measure works based on two compounds.

Each patent includes numerous compounds[7] and to calculate their similarity, they can be compared using a clustering technique. For this step, we leverage a clustering approach called binning clustering, using the ChemMine software[8]. Binning clustering assigns each compound to a group of similar compounds based on a similarity cut-off, typically close to 0.4 (Muthas and Boyer, 2013). Using this clustering technique, we assess the similarity of pairs of patents based on the clustered compounds. If two patents have few (many) clustered compounds in common, they are considered dissimilar (similar) in the chemical space based on their chemical characteristics. In other words, the more the similarity between two patents, the more the overlap of their clustered compounds. To calculate this overlap, we rely on a standard Jaccard index (Nikolova and Jaworska, 2003):

[6] Krieger et al. (2022) use a very similar approach to assess the degree of novelty of new molecules submitted to trials by pharmaceutical firms.

[7] Each patent included in our set reports on average 122.42 compounds (Std. dev. 181.57).

[8] See https://chemminetools.ucr.edu/.

$$J_{A,B} \equiv \frac{|A \cap B|}{|A \cup B|} = \frac{|A \cap B|}{|A| + |B| - |A \cap B|} \tag{1}$$

This measure captures the fraction of chemical features that are shared by the two patents and takes values between 0 and 1, with 0 implying no chemical features in common. Figure 3 shows a simplified example of how this measure works by comparing pairs of patents using the Jaccard index. In this case, the two patents have a similarity value of 0.55, resulting from 5 compounds clustered together.

Figure 3: Example of the similarity measure between two patents

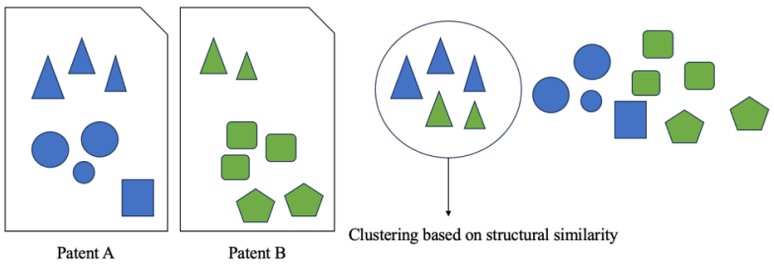

Jaccard=5/(7+7-5)=5/9=0.55

Following the approach by Krieger et al. (2022), we then compute the within firm similarity for each year. For each firm i at time t, we compute the Similarity$_{i,t}$ as patents filed at time t maximum pairwise similarity to patents filed in the previous three available years.[9]

Figure 4: Distribution of similarity

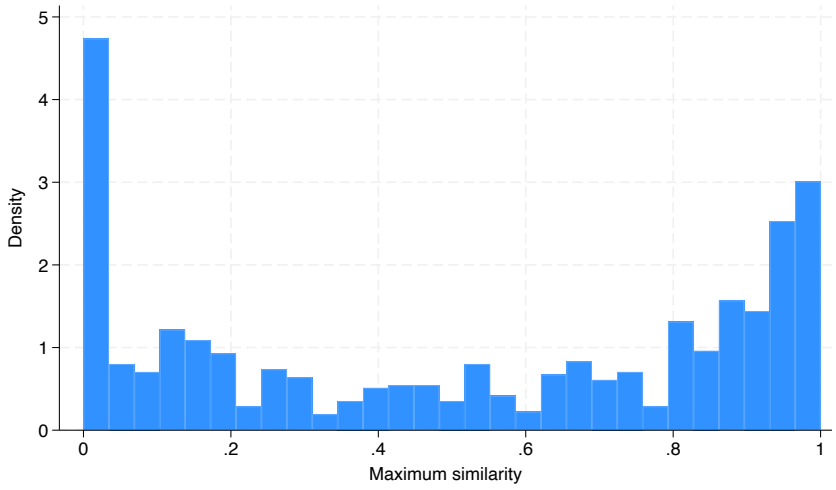

Distribution of within firms similarity in patents filed in the domain of POPs.

Figure 4 shows the within-firm similarity measure distribution. The graph shows that the distribution is bimodal, with most patents having maximum similarity close to 0 or close to 1. Recalling that a low similarity score implies high novelty, this evidence suggests that most patents are either very novel or not

[9]See Appendix C for an example of how we compute the firm's maximum similarity.

novel. Krieger et al. (2022), the only previous paper using a very similar metric, also finds this bimodal distribution of patent novelty in the pharmaceutical industry.

Figure 5: Patent similarity between small and large firms

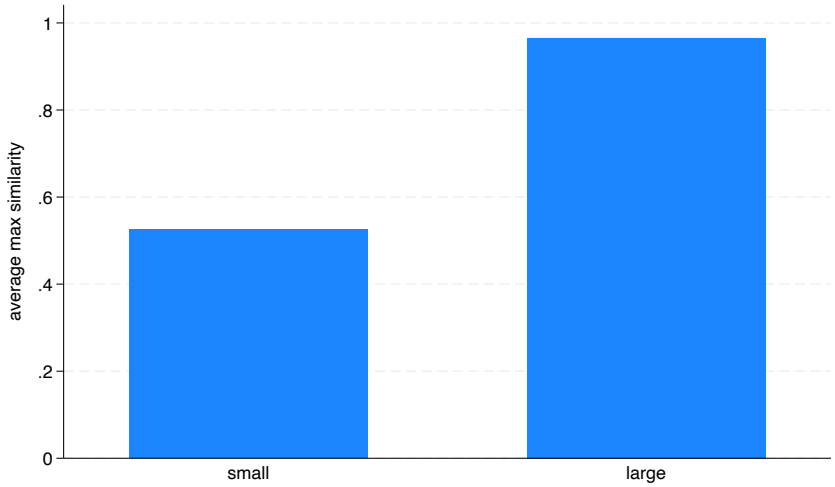

Within firms similarity in patents filed in the domain of POPs between small and large firms. Small and large firms are defined by the firm's patent portfolio in each country. Large firms have a patent portfolio above the 95th percentile of the distribution. Small firms belong to the residual category.

Figure 5 displays how the average degree of novelty differs between large and small firms regarding their patent portfolio. We define large firms as those with a patent portfolio above the 95th percentile of the portfolio's distribution. Interestingly, larger firms file patents with a lower degree of novelty compared to small firms.

Finally, future effort will be dedicated to build a third outcome variable capturing the degree of toxicity of the compounds included in the patented invention (Biggi et al., 2022). Given the presence of patents with high level of compound structural similarity, we expect that the new patents are still hazardous for human health and environment even if not strictly banned. Using together the similarity and toxicity measure we aim testing whether the bans has moved firms' research efforts towards the development of new chemical structure possibly less toxic.

4.2 Difference-in-differences with multiple time periods

Our empirical strategy follows the estimation procedure of Callaway and Sant'Anna (2021). This approach accounts for variation in treatment timing and heterogeneous treatment effects, in the presence of which the standard two-way fixed effects difference-in-difference model does not guarantee to estimate an interpretable causal parameter (Borusyak and Jaravel, 2018; De Chaisemartin and d'Haultfoeuille, 2020; Goodman-Bacon, 2021). In our approach, the control group is composed of the not-yet treated countries, excluding the never treated (i.e. those few countries that have never experienced a ban).[10]

[10] Our analysis is performed using the CSDID package in STATA (Rios-Avila et al., 2023).

At this stage, we limit the period of analysis between 1960 and 2000 to avoid possible confounding factors stemming from the Stockholm Convention and the potential influences this global ban may have had on firms' innovation activity.

As the primary dependent variable, we use the logarithm of the number of POPs-related patents filed by the firm i in country k in year t. As control variables, we include the number of jurisdictions in which the firm is patenting and a dummy equal to one if the applicant is a university or research centre.

5 Results

Table 3 displays the results of the regression analysis to test our first hypothesis. In the first column, we use the count of POPs-related patents filed by each firm i in the country k in log terms as an outcome variable. In the second and third columns, we distinguish between chemical and non-chemical POPs-related patents using the classification developed by Van Looy et al. (2015).[11] In all specifications, we include the number of jurisdictions in which the firm is patenting as control variables and a dummy if the applicant is a university or a research centre. Moreover, firm- and year- fixed effects are included. We cluster the standard error at the firm-country level.

Table 3: Firm-level staggered diff-diff results

	all POPs	POPs applied in chemical	POPs applied outside chemical
ATT	-0.00894	-0.0254**	0.0128*
	(0.010)	(0.011)	(0.007)
Observations	15,968	15,968	15,968
Year FE	✓	✓	✓
Firm-country FE	✓	✓	✓
Additional controls	✓	✓	✓
Period of analysis: 1962-2000.			
Standard errors in parentheses, clustered by firm-country.			
Legend: *** p<0.01, ** p<0.05, * p<0.1			

Our findings suggest that country-level initiatives aimed at banning POPs do not affect firms' innovation activity related to POPs. When using the total number of POPs-related patents as the outcome variable, the Average Treatment Effect of the Treated (ATT) is not significant. However, by distinguishing POPs applied to chemical or non-chemical fields, we get an interesting insight: bans have a negative impact once POPs are applied to chemical technologies, while they have a positive (but barely significant) effect on POPs patents applied outside chemical fields. A graphical illustration of our results in time is shown in Figure 6. This result suggests that firms are possibly "repurposing" POP-related inventions into technological domains not directly associated with their primary chemical use. Overall, these bans

[11] This classification is based on the IPC class of each patent, and it is retrieved from the table TLS229 of the EPO-PATSTAT database.

have mixed effects on POP-related innovation activities, only partially confirming our first hypothesis. Investigating the effect on similarity and toxicity should help disentangle which channels are affected mainly by the regulatory change.

Figure 6: Effect of the bans on POPs-related patents applied in chemical and non-chemical fields

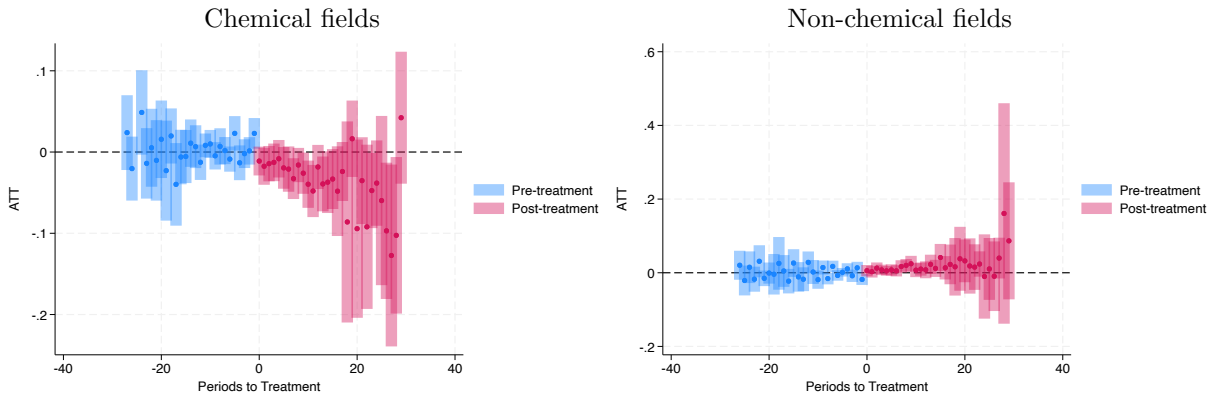

These results are robust if we exclude universities and research centers from the sample. Moreover, they are also robust if we run the regression at the country-level, as shown in Table A3.

Figure 7: Average max similarity across countries in time

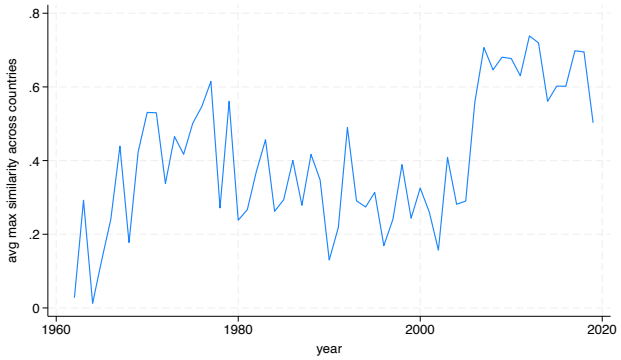

The max similarity measure is computed at the patent-level.

Moving to our second hypothesis, we do not yet have formal testing using regression analysis, so we provide some preliminary descriptive results. Figure 7 displays the average maximum similarity across countries over time. In this case, the similarity measure is calculated at the patent level, computing the maximum similarity between each patent filed in a year and all patents previously filed in the same country. Following the second hypothesis, we expect that bans decrease the similarity to POP patents, indicating a different innovative trajectory over time. However, figure 7 does not display such a decline. Instead, on average, patents seem to become more similar between 1960 and 1980 and from 2000. This descriptive evidence suggests that firms might not change their innovative trajectory following the bans but continue on the same trajectory, patenting compounds similar to POPs but possibly not banned. The issue with very similar compounds is that structural similarity entails similar characteristics, such

as toxicity and potential harm to human health and the environment. Despite preliminary, this evidence goes against our second hypothesis. Future efforts will be dedicated to testing the second hypothesis using the regression approach discussed in Section 4.2.

6 Conclusion

Evidence of the dangers of POPs was already available in the '60. Since 1962, several countries have started introducing regulations intending to phase out the production, use, import, and export of POPs. While country-specific bans were designed to forbid future production and use of these chemicals, they did not directly regulate the R&D processes underpinning the banned compounds, perhaps based on the presumption that these would have been phased out too. Our study reveals that country-level initiatives aimed at banning POPs effectively reduced patenting activities in the domain of POPs only in chemical fields. In contrast, the number of POP applications filed outside the chemical fields has increased after the ban. Investigating the effect on similarity and toxicity should help disentangle which channels are affected mainly by the regulatory change. Preliminary evidence suggests that similarity does not decrease over time, highlighting how firms might introduce new compounds that are not banned but very similar to POP and, therefore, very toxic.

References

Bajusz, D., A. Rácz, and K. Héberger (2015): "Why is Tanimoto index an appropriate choice for fingerprint-based similarity calculations?" *Journal of cheminformatics*, 7, 1–13.

Bartrons, M., J. Catalan, and J. Penuelas (2016): "Spatial and temporal trends of organic pollutants in vegetation from remote and rural areas," *Scientific reports*, 6, 1–10.

Biggi, G., E. Giuliani, A. Martinelli, and E. Benfenati (2022): "Patent toxicity," *Research Policy*, 51, 104329.

Borusyak, K. and X. Jaravel (2018): *Revisiting event study designs*, SSRN.

Brusoni, S., A. Prencipe, and K. Pavitt (2001): "Knowledge specialization, organizational coupling, and the boundaries of the firm: why do firms know more than they make?" *Administrative science quarterly*, 46, 597–621.

Callaway, B. and P. H. Sant'Anna (2021): "Difference-in-differences with multiple time periods," *Journal of Econometrics*, 225, 200–230.

Carson, R. (2002): *Silent spring*, Houghton Mifflin Harcourt.

Cousins, I. T., J. H. Johansson, M. E. Salter, B. Sha, and M. Scheringer (2022): "Outside the safe operating space of a new planetary boundary for per-and polyfluoroalkyl substances (PFAS)," *Environmental Science & Technology*, 56, 11172–11179.

De Chaisemartin, C. and X. d'Haultfoeuille (2020): "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 110, 2964–2996.

de Rassenfosse, G., W. E. Griffiths, A. B. Jaffe, and E. Webster (2021): "Low-quality patents in the eye of the beholder: Evidence from multiple examiners," *The Journal of Law, Economics, and Organization*, 37, 607–636.

Dosi, G. (1982): "Technological paradigms and technological trajectories: a suggested interpretation of the determinants and directions of technical change," *Research policy*, 11, 147–162.

Goodman-Bacon, A. (2021): "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 225, 254–277.

Jepson, P. D. and R. J. Law (2016): "Persistent pollutants, persistent threats," *Science*, 352, 1388–1389.

Johnson, M. A., G. M. Maggiora, et al. (1990): "Concepts and applications of molecular similarity," *(No Title)*.

Krieger, J., D. Li, and D. Papanikolaou (2022): "Missing novelty in drug development," *The Review of Financial Studies*, 35, 636–679.

Ma, J., H. Hung, C. Tian, and R. Kallenborn (2011): "Revolatilization of persistent organic pollutants in the Arctic induced by climate change," *Nature Climate Change*, 1, 255–260.

Magerman, T., B. Van Looy, and X. Song (2006): "Data production methods for harmonized patent statistics: Patentee name harmonization," .

March, J. G. (1991): "Exploration and exploitation in organizational learning," *Organization science*, 2, 71–87.

Muthas, D. and S. Boyer (2013): "Exploiting pharmacological similarity to identify safety concerns–listen to what the data tells you," *Molecular Informatics*, 32, 37–45.

Nikolova, N. and J. Jaworska (2003): "Approaches to measure chemical similarity–a review," *QSAR & Combinatorial Science*, 22, 1006–1026.

Nizzetto, L., M. Macleod, K. Borgå, A. Cabrerizo, J. Dachs, A. Di Guardo, D. Ghirardello, K. M. Hansen, A. Jarvis, A. Lindroth, et al. (2010): "Past, Present, and Future Controls on Levels of Persistent Organic Pollutants in the Global Environment," *Environmental Science & Technology*, 17, 6526–6531.

Özkara, A., D. Akyil, and M. Konuk (2016): "Pesticides, environmental pollution, and health," in *Environmental health risk-hazardous factors to living species*, IntechOpen.

Rigét, F., A. Bignert, B. Braune, M. Dam, R. Dietz, M. Evans, N. Green, H. Gunnlaugsdóttir, K. S. Hoydal, J. Kucklick, et al. (2019): "Temporal trends of persistent organic pollutants in Arctic marine and freshwater biota," *Science of the total environment*, 649, 99–110.

Rios-Avila, F., P. Sant'Anna, and B. Callaway (2023): "CSDID: Stata module for the estimation of Difference-in-Difference models with multiple time periods," .

Rockström, J., W. Steffen, K. Noone, Å. Persson, F. S. Chapin, E. F. Lambin, T. M. Lenton, M. Scheffer, C. Folke, H. J. Schellnhuber, et al. (2009): "A safe operating space for humanity," *nature*, 461, 472–475.

Van Looy, B., C. Vereyen, and U. Schmoch (2015): "Patent Statistics: Concordance IPC V8-NACE Rev. 2 (version 2.0)," *Final report, EUROSTAT, Luxembourg.*

A Tables and figures

Table A1: Datasources

Datasource	Type of data
CAS SciFinderN	chemical structure and its functions in the product/process sought by the patented invention
EPO - PATSTAT	patent-related information (e.g. date of filing, applicant, IPC class).
Own dataset	hand-collected dataset of bans across countries
ORBIS - IP	Firm-level information (e.g. financial data)

Table A2: Dataset composition

Period of analysis	1960-2016
Nr of:	
regulating countries	146
for which:	
patent data are available	85
patents with POPs compounds	3,555
among which:	
filed in chemical IPC classes	1,844
filed in non-chemical IPC classes	1,381
info not available	330
among which:	
filed in agro fields	1,852
filed in non-agro fields	1,703
patent applicants	898
among which:	
private firms	695
universities/institutions	203

Notes: Agro-fields are defined by the section title available in SciFinderN.

Figure A1: Evolution of POPs-related patent applications by country

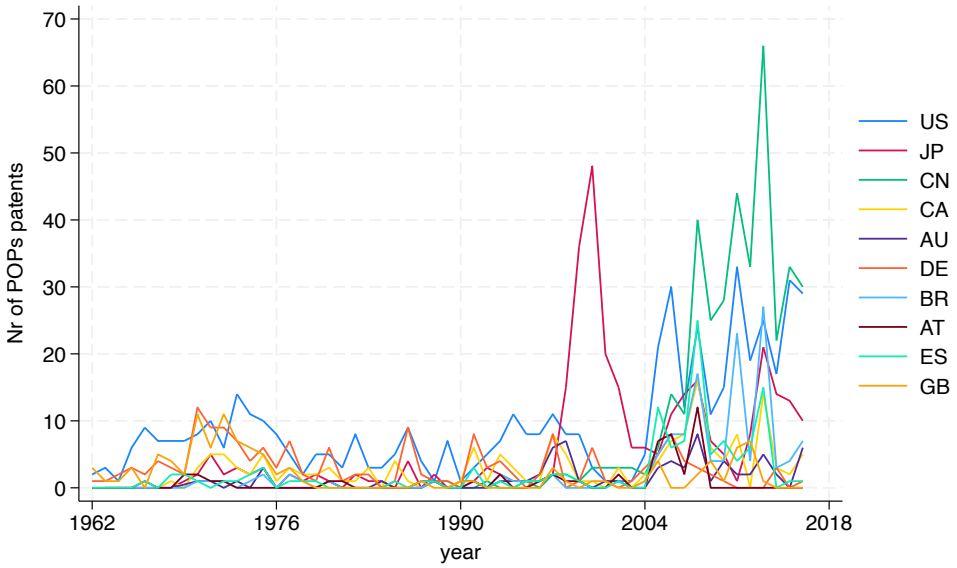

Our own elaboration on the top patent offices worldwide in terms of number of POPs related patent applications filed between 1962 and 2016. The displayed 10 patent offices cover the 50% of the patents in our dataset.

Table A3: Country-level staggered diff-diff results

	all POPs	POPs applied in chemical	POPs applied outside chemical
ATT	-0.112	-0.186*	0.0307
	(0.098)	(0.099)	(0.082)
Observations	1,482	1,482	1,482
Year FE	✓	✓	✓
Country FE	✓	✓	✓
Period of analysis: 1962-2000.			
Standard errors in parentheses, clustered by country.			
Legend: *** p<0.01, ** p<0.05, * p<0.1			

B Computing the Tanimoto coefficient

Figure B1 displays two similar alcohol compounds, ethanol (Compound A) and methanol (Compound B). Using the Tanimoto coefficient formula (Equation 1), we can calculate their structural similarity based on the presence of common chemical features (fragments). In this case, we have two common features at the intersection (CH3 and OH) and three unique features at the union (CH3, CH2, OH), resulting in a Tanimoto coefficient of approximately 0.67. This measure suggests that ethanol and methanol share two of three selected structural features, making them similar.

Figure B1: Example of two structurally similar compounds

Compound A: Ethanol

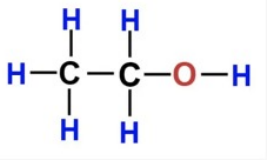

Chemical Structure: C2H5OH
Simplified Representation: CH3-CH2-OH

Compound B: Methanol

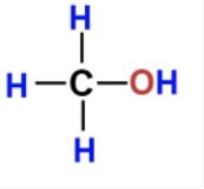

Chemical Structure: CH3OH
Simplified Representation: CH3-OH

C Computing firm similarity

Figure C1 displays a numerical example of how the within firm similarity is calculated.

Figure C1: Example of firm similarity