International Environmental Agreements and Directed Technological Change: Evidence from the Ozone Regime*

Eugenie Dugoua†

July 2, 2018

Abstract

Can international environmental agreements induce innovation on green technologies? It is possible that international negotiations succeed only once technological solutions are available. In this case, agreements would help diffuse such technologies rather than fostering their development. I provide the first quantitative evidence that the Montreal Protocol, and its following amendments to protect the ozone layer, triggered a large increase in research and innovation on alternatives to ozone-depleting molecules. To do this, I use the full text of patents and scientific articles to construct new panel data of the yearly number of published documents about these molecules. I implement a difference-in-differences strategy (DiD) and a synthetic control method (SCM) using hazardous air pollutants as control units. To compare molecules' chemical and industrial characteristics, I construct descriptive variables by applying machine learning techniques to the documents' text. The SCM estimates that the post-Montreal regime caused a 144% increase in patents and a 189% increase in articles mentioning substitutes to ozone-depleting substances; the DiD yields comparable estimates. These results challenge the view that agreements foster technological diffusion without affecting much of the dynamics of innovation. Agreements can thus encourage the development of green technologies, which importantly suggests they should be negotiated as early as possible if we hope to solve global environmental problems.

JEL Classification: O31, O33, Q55, Q58, F53
Keywords: ozone, CFCs, Montreal Protocol, international environmental agreements, technological change, difference-in-differences, synthetic control, topic modeling

*I am deeply grateful to Scott Barrett, Suresh Naidu, Geoffrey Heal for their advice and support. I thank William McAllister, Cristian Pop-Eleches, Bernard Salanie, Johannes Urpelainen, Doug Almond, Wolfram Schlenker, and Rodriguo Soarez for their advice. I also thank Marion Dumas, Anna Thompsett, Kyle Meng, Prabhat Barnwal, Jeffrey Shrader, Anthony D'Agostino, Eyal Frank, Anouch Missirian and Johannes Boehm for their helpful comments, as well as the participants of the Applied Microeconomics Workshop, the Sustainable Development Seminar and Colloquium, the Interdisciplinary PhD Workshop on Sustainable Development, and all the fellows of the Mellon Foundation. All errors are mine.

†Assistant Professor in Environmental Economics, London School of Economics (starting Sept.). E-mail: eugenie.dugoua@columbia.edu.

Contents

1 Introduction

International environmental agreements, like domestic policies, attempt to mitigate environmental degradation. Agreements, however, are needed when environmental problems run across national borders. In such contexts, individual governments usually lack incentives to create domestic policies, and international cooperation is needed. A large literature has developed to better understand the drivers of environmental-friendly innovation, and many studies show that domestic environmental policies can foster green innovation (Jaffe et al. 2002; Popp et al. 2010). I ask: Can international environmental agreements, like domestic policies, foster innovation? It is possible that they don't: noncooperative game theory suggests that agreements occur when costs to the players are low. Hence, agreements might only occur once technological solutions are readily available, and they might simply contribute to diffusing these technologies, as opposed to fostering the development of new ones.

This paper shows that, on the contrary, agreements can induce innovation and that agreements, therefore, are part of the process of delivering cheaper environmental-friendly technologies. This provides a strong argument for negotiating ambitious agreements as early as possible since technologies are too often the keystone for addressing environmental problems. To make this argument, I provide empirical evidence from the Montreal Protocol and its following amendments. In 1987, at Montreal, high-income countries decided to phase-out chlorofluorocarbons (CFCs) from industrial activities because CFCs were known to destroy the protective layer of ozone molecules in the stratosphere. Technological change unrolled rapidly: within a decade, the production and consumption of CFCs decreased by more than 80%[1]. The protocol is still hailed as one the most successful environmental international agreement and remains a point of reference in policy discussions about global environmental problems.

Despite the large scholarly literature on the topic, the dynamics of innovation in the ozone crisis are still debated. Richard E. Benedick, chief U.S. negotiator at Montreal, claims the agreement triggered a vast effort in research to find CFC substitutes (Benedick 2009). But others emphasize that CFC substitutes were already available at the time of negotiations (Heal 2016; Sunstein 2007).

[1] My calculations using UNEP data.

This paper is not only the first to empirically show that Montreal fostered innovation, but it also quantifies its effect. I do this using a novel molecule-level panel dataset with both a difference-in-differences (DiD) strategy and a synthetic control method (SCM). Additionally, I apply machine learning methods to semantically match documents and measure similarity between molecules.

Successfully developing CFC substitutes did not really rely on identifying chemical structures because the set of molecules with greatest potential to be CFC substitute was already well-known at the time; I compile a list of 14 of such molecules[2] and consider those molecules as treated by the Montreal Protocol. The technological challenge, instead, lied in finding out how such molecules could be used in the myriad of industrial processes that required CFCs, cost-effectively and at a large scale. This meant, first, learning about thermodynamics properties, toxicity profile, and environmental acceptability; I collect scientific articles published in journals indexed by Science Direct between 1970 and 2000 to capture such research effort. Second, new process and formula designs were needed to retrofit installed equipment with the CFC substitutes or to altogether replace it; I collect patents granted by the United States Patent and Trade Office (USPTO) between 1976 and 2000 to track progress on these aspects.

Unfortunately, no preexisting classifications allow me to easily identify which documents relate to CFC substitutes. I, therefore, search the full text of patents and articles for mentions of any of the 14 CFC substitutes[3] and construct a panel dataset where each observation is the number of documents mentioning a given molecule at least once in a given year. Finally, I apply machine techniques to text analysis to construct variables that proxy the scientific and industrial context of the molecules. These variables correspond to the proportion of specific topics present in documents mentioning molecule i. Intuitively, they describe the type of words associated with molecule i.

I begin by estimating the difference before and after the signing of Montreal in the number of documents mentioning CFC substitutes: I find large increases, close to 600% in patents and 200% in articles. Additionally, I find that only very few patents and articles on CFC substitutes

[2]I used the report published in 1988 by the AFEAS (Alternative Fluorocarbon Environmental Acceptability Study). After the agreement at Montreal, manufacturers were given authorization by anti-thrust authorities to cooperate on some specific areas. The AFEAS publication shared what they knew about the *atmospheric characteristics* of several potential CFC substitutes.

[3]Since molecules usually have many different names[4], I develop an automatic script to collect all possible names from SciFinder, a database of chemical information.

are published before 1987, and the trend prior to 1987 ("pre-trend") is remarkably flat. I argue that this is indicative of technological progress not being a key driver of negotiations' success, and that the agreement was little anticipated. To account for potential underlying trends, I compare innovation on CFC substitutes with a control group: I use 171 hazardous air pollutants (HAPs). These molecules can serve as controls because they are unrelated to ozone or to CFCs and, just like CFC substitutes, are used in a diverse range of industrial applications. Importantly, pre-trends in the number of documents mentioning both sets of molecules are comparable. The main DiD estimate indicates that Montreal led to an increase of about 546% and 95% in the number of patents and articles, respectively, between 1987 to 2000. This corresponds to average annual increases of about 390 patents and 47 articles. The estimates are reduced but remain economically and statistically significant when controlling for lags and topic proportions.

Since one patent or article can often mention several CFC substitutes, the observations used in the DiD design (14 CFC substitutes) are not independent. Another approach, therefore, consists in considering them in aggregate, as one treated molecule. To estimate a treatment effect on such "aggregate CFC substitute", I use an SCM (Abadie et al. 2010, 2015), a method particularly suited for estimating treatment effect of interventions affecting aggregate quantities. To do this, the SCM constructs a control unit by using a weighted average of control molecules. The method chooses the weights so that the synthetic control unit reproduces most closely the log count path of the outcome variable in the pretreatment periods. I also use weights chosen so that the synthetic control resembles the treated unit along topic proportions; this helps obtain a control unit that resembles CFC substitutes along chemical and industrial dimensions.

The average treatment effect using the SCM indicates the Montreal protocol lead to an increase of about 144% in the number of patents mentioning the CFC substitutes, corresponding to about 117 additional patents per year over the study time period after 1987. This yields a lower estimate than the DiD strategy, indicating that the control constructed in SCM provides a more conservative comparison. For articles, on the other hand, results are similar to the DiD. I find an average treatment effect close to a 190% increase in the number of articles, which corresponds to about 43 additional articles per year after 1987. To assess the statistical significance of these results, I

follow the placebo tests method suggested by Abadie et al. (2010, 2015). I find treatment effects are significant at the 99% level. In addition, the increase in the number of documents mentioning CFC substitutes becomes statistically significant as from 1990, three years after the agreement was signed. This lag might correspond to the time it organizationally and professionally takes to redirect research towards CFC substitutes and to have patents and articles published (Popp 2002).

These results support the idea that the post-Montreal ozone regime caused the development of CFC substitutes. This finding is robust to considering citation-weighted document counts; indeed, the most cited patents and articles on CFC substitutes were published after Montreal. Additionally, one hypothesis is that manufacturers kept their CFC substitutes secret, and Montreal simply created a world market for them. If that were true, we would observe a sudden increase in patents in the few months following the treaty signature. I show that this is not the case both for all patent assignees as well as for the biggest two, DuPont and Dow Chemical. Additionally, I find that these results are robust to dropping patents and articles with only few mentions of CFC substitutes (e.g. keeping documents with at least three occurrences of a molecule name).

This paper contributes to the literature on directed technological change in the context of environmental issues (Acemoglu 2002; Jaffe et al. 2002; Popp 2010). Many papers have explored the relationship between domestic environmental regulations and innovation (e.g., Aghion et al. (2016), Calel et al. (2016), and Jaffe et al. (1997)). My paper is most similar to Dekker et al. (2012) in that it investigates the causal effect of an international agreement[5]. However, I focus on a case of a global public good (strastospheric ozone). Importantly, while scholars have thoroughly investigated the diplomatic and game theoretic aspects of the ozone crisis (Barrett 2003; Benedick 2009; Murdoch et al. 2009; Parson 2003; Wagner 2009), no quantitative analysis of the dynamics of innovation during the crisis has been carried out. This is despite the economics, science, and politics of ozone serving as an anchor point for our understanding and beliefs about the role of diplomacy, agreements, and technologies in solving environmental issues, especially climate change (Barrett 1999; Sunstein 2007). This paper thus complements the literature on Montreal by showing

[5]Dekker et al. (2012) focus on the signing of the Helsinki and Oslo protocol which aimed at reducing transboundary sulfur emissions

and quantifying its effect on science and innovation. When solutions to environmental problems are plagued with technological uncertainties or high price tags, decision-makers are incentivized to adopt a "wait-and-see" strategy: wait for proven new technologies, then negotiate an agreement. By showing that agreements can encourage the development of green technologies, this paper suggests they should be negotiated as early as possible if we hope to solve global environmental problems.

The following section 2 summarizes the literature on directed technological change and its relationship to the environment and provides further information on the ozone crisis and the Montreal Protocol. I describe the data in section 3 and the empirical strategies and main results in section 4 and 5. I discuss some caveats in section 7 and conclude in section 8.

2 Directed Technological Change and the Ozone Layer

2.1 Directed Technological Change and the Environment

The relationship between technological change and the environment has been drawing more interest, particularly since the 1990s. On the one hand, technical change affects the intensity of environmental impacts. On the other, there is the growing recognition that environmental policies create new types of incentives and constraints possibly affecting the direction of technological change. The concept of *directed technological change* goes back to 1936: under the "induced innovation" hypothesis, Hicks (1932) stated that innovations are biased towards high priced factors so to make their use more efficient or to substitute them. In the past two decades, the concept has reappeared under the phrase "directed technical change" (Acemoglu 1998) encompassing not just price effects, but also market size and regulatory effects[6].

In the environmental context, the direction of innovation is particularly important. The usual technology policy (e.g., public funding for research and development activities or intellectual prop-

[6]The phrase's popularity took off after the publication of an article by Daron Acemoglu (Acemoglu 1998) showing that an increase in skilled labor force can induce skill-biased technological change through a market size effect that fosters the development of innovations complementary to the abundant factor, in that case, skills. In another paper, Acemoglu (2002) presents a model where the direction of technological change is influenced by both scarce factors (through prices) and abundant factors (through market size).

erty regimes) attempts to deal with knowledge market failures by fostering the rate of innovation and diffusion of new technologies. But it does so in a direction-blind way. As a result, a large literature has developed at the intersection of environmental and technology policy to better understand how and to what extent technical change could be directed (Jaffe et al. 2002; Popp 2010; Popp et al. 2010). In fact, we can think of environmental regulations as modifying the shadow prices of environmental inputs which, as the induced innovation hypothesis suggests, induces innovation in non-polluting directions. This is specifically discussed by Newell et al. (1999) who generalized the concept of induced innovation to include inducement by regulations.

The literature initially focused on the impact of environmental regulations on business competitiveness (Ambec et al. 2013; Porter 1991; Porter et al. 1995a,b) and then later on patenting activities and R&D spendings; scholars found strong evidence that regulations have an important influence on environmental-friendly innovations (Brunel 2015; Brunnermeier et al. 2003; Jaffe et al. 1997; Johnstone et al. 2010, 2012; Lanjouw et al. 1996; Nesta et al. 2014; Popp 2005; Vollebergh 2007). For example, Popp (2006) finds significant increases in patents pertaining to sulfur dioxide and nitrogen oxides emissions reduction in response to the passage of environmental regulations in the United States, Japan, and Germany. More recently, Calel et al. (2016) show that the European Union Emissions Trading System increased patenting related to low-carbon technologies by about 10%, while not crowding out other technologies. Although many studies investigate the effect of international environmental agreements on pollution outcomes (Aichele et al. 2011; Finus et al. 2003; Kellenberg et al. 2014), they seldom look at the impact on science and innovation. One exception is Dekker et al. (2012) who show increased patenting activity for countries signatories of the Convention on Long-Range Transboundary Air Pollution.

2.2 A Brief History of the Ozone Crisis

The story of the ozone crisis began in 1974 when two chemists published an article in which they laid out the theoretical possibility that ozone molecules could be broken down in the stratosphere by chlorofluorocarbons (CFCs) (Molina et al. 1974)[7]. Even though the potentially harmful effects

[7]In 1995, Mario J. Molina and F. Sherwood Rowland, together with Paul J. Crutzen, were awarded the Nobel Prize in Chemistry "for their work in atmospheric chemistry, particularly concerning the formation and decompo-

of a thinner ozone layer were not well understood, it was clear that more UV light would cause more skin cancers, eye cataracts and a likely loss in fishery and agriculture productivity (Miller et al. 1986). CFCs, the main molecules responsible for depleting ozone, had become important molecules for industrial activities due to their chemical properties: they are unusually stable, nonflammable, nontoxic and noncorrosive[8]. This makes them ideal molecules for manufacturing many consumer goods. And, best of all, they were cheap to produce. The use of CFCs spread over mostly five different sectors: foams, refrigeration and air-conditioning, aerosols, fire protection and solvents[9].

In September 1987, industrialized countries agreed to a binding agreement regulating the production and use of CFCs. The approach was flexible with a series of phase-out dates, as opposed to banning altogether the molecules. These phase-out schedules were further consolidated and extended to other molecules in the years that followed[10]. Atmospheric concentrations for most CFCs peaked by 2014 and ozone layer recovery is now expected around 2050 (Hegglin et al. 2015). The role of the Montreal Protocol in solving the crisis has been intensely discussed (Barrett 1994; Beron et al. 2003; Murdoch et al. 1997; Wagner 2009, 2016). Specifically, Barrett (Barrett 1999) suggested that a key aspect of the protocol was to solve the enforcement problem: Montreal included trade restrictions with non-parties in ozone-depleting substances as well as in products containing those substances. It also included the threat of banning trade in products made using ozone-depleting substances. These trade restrictions effectively acted as a mechanism for free-rider deterrence and leakage prevention. More recently, Wagner (Wagner 2016) provided empirical evidence that these trade measures promoted full participation in the protocol, ensuring its almost-universal ratification. My paper complements this literature by focusing on the

sition of ozone"

[8]Initially, CFCs somewhat embodied the miracle of modern chemistry. They were first commercially used in 1928 as cooling fluids for refrigerators, and were specifically designed to substitute other dangerous refrigerants that were either toxic or inflammable (Parson 2003).

[9]CFCs are great refrigerants because they vaporize at low temperature and are very energy-efficient coolants. As aerosols, they were used in cosmetics, household products, pharmaceuticals, and cleaners. Finally, their nonreactive property made key products for cleaning microchips and telecommunication equipment

[10]For example, the London amendment, signed in 1990, regulated new chemicals such as carbon tetrachloride and methyl chloroform. In 1995, the parties successfully negotiated phase-out targets for lower-income countries, which were until then exempted from any regulation.See supporting online material for a detailed schedule for the targets.

quantitative effect of Montreal and its following agreements on science and innovation.

2.3 The Role of Technology in the Ozone Crisis

There is no question that if emission reductions were successful, it was thanks to CFC substitutes becoming available. Goods that contained or required CFCs for their production continued to be commercialized, and no air capture system of CFCs was ever designed. But the question of when these CFC substitutes were developed and whether the agreement triggered the bulk of the effort to find them is still debated. Some works have focused specifically on the technological story behind the Montreal Protocol (Glynn 2002; Gonzalez et al. 2015; Le Prestre et al. 1998; Miller et al. 1986; Parson 2003; Taddonio et al. 2012) as well as the reaction of the business community (Falkner 2005; Mulder 2005; Reinhardt et al. 1989a,b; Smith 1998). But, perspectives on the role of innovation remain mixed.

Richard Benedick, head U.S. negotiator at Montreal, argued that the agreement caused a vast effort in research to find CFC substitutes and that qualitative evidence abounds on the dynamics of the innovation process under the Montreal Protocol (Benedick 2009)[11]. Similarly, Edward A. Parson highlights that, although some manufacturers initially started research on potential substitutes in the late 1970s, these efforts quickly came to an end around 1981 (Parson 2003, Chap.3 p.53 and Chap.7 p.173)[12]. As a result, little was known about the toxicity and environmental acceptability of the potential CFC substitutes, whether and how they could be processed at large scale, and whether they would require a redesign of the processes and equipment in the various industries that used CFCs as inputs. Despite both Benedick and Parson arguing that, on the eve of the negotiations, technological uncertainties loomed large, others have taken a different stance. One narrative claims that CFC substitutes were readily available before the

[11] Benedick (2009, Chap.8 p.104.): "It was evident (...) that the protocol was in fact moving industry in directions that two years earlier had been considered impossible." Benedick refers to articles published in the *New York Times* and *Chemical and Engineering News* when he asserts that the agreement triggered a vast research effort.

[12] Parson narrates the various waves of research efforts to develop CFC substitutes, both before and after the signature of the Montreal Protocol. According to Parson, manufacturers would have stopped these R&D programs because they had determined that CFC substitutes would cost around two to five times more than CFCs, and it made no sense to continue developing these substitutes with little sign of regulations under way.

negotiations (Heal 2016; Sunstein 2007)[13]. This view is also often expressed in media outlets[14]. In this paper, I empirically investigate the role of the Montreal protocol and its following amendment in the development of CFC substitutes by quantitatively analyzing trends in patents and articles related to CFC substitutes.

3 Data

3.1 Empirical Indicators of Technological Change

Patents. The process of technological change is often described as a sequence of distinct activities: basic research, applied research, development, commercialization, and diffusion (Greenhalgh et al. 2010, figure 1.1). To strengthen the incentives to generate innovations, modern economies have adopted regimes of intellectual property rights where inventors are granted exclusive rights through patents. The publication of a patent, as a result, is a testimony to the successful process of applied research and development and can be considered as a proxy for technological change in general, and of innovation in particular. Patent data has been broadly used in empirical research in the past two decades (Hall et al. 2012; Henderson et al. 1998; Kay et al. 2014; Popp 2005; Williams 2013; Williams 2017). I follow this literature by using patent counts as a proxy for innovation. I download the full-text of all U.S. patent grants from 1976 to 2000 from the U.S.P.T.O. repository[15]. To construct a better proxy for innovation, I sort patent by application date. I use the texts contained in the abstract and summary description of the invention[16]. The

[13]Sunstein (2007), for example, claims that "an international agreement was largely in the interest of American manufacturers, which had already initiated a transition to safe CFC-alternatives". In a recent book examining the most urgent environmental issues of our time, Geoffrey Heal discusses the failure of the Kyoto Protocol in comparison to the success of Montreal: "However, there are big differences between ozone depletion and climate change. We have not yet seen the equivalent of DuPont's discovery of an alternative to CFCs, which would be the discovery by oil and coal companies of a greenhouse-gas-free energy source capable of meeting world energy demand at current energy costs." (Heal 2016, p68).

[14]Here is an excerpt from an article published in *The New York Times* on August 20, 2002: "The agreement's success occurred, in large part, because substitutes for the harmful chemicals were readily available (...)."

[15]This represents a total of 2,605,925 patents.

[16]Patents include an abstract, a description of prior art, a summary description of the invention, a detailed description and a list of claims. Since the writing of patents is subjected to a level of scrutiny much higher than for scientific articles, abstracts and summary descriptions are likely, in this case, to faithfully represent the invention. I, therefore, limit the study to those sections only. Reproducing the whole analysis using the text of the detailed description of the invention and all the claims would be possible, at the expense of computational time.

cleaning procedure involves a series of standard steps such as replacing English contractions with their non-shortened forms or converting non-ASCII characters into their closest ASCII equivalents. More information is provided in Appendix A. Patents also contain the names and addresses of inventors and assignees[17]. I also categorize patents by the type of organizations their assignee is affiliated with (e.g., business, education, or government). More details about how the meta-data is cleaned, matched and classified by type are provided in Appendix A. For patent citations, I use the NBER U.S. Patent Citations Data File[18].

Articles. The phases of basic and applied research focus more specifically on the production and dissemination of knowledge. Through the publication of articles, researchers render their contribution public and allow other researchers and potential inventors to build upon it. A large literature was born out of Derek J. de Solla Price's contributions to a quantitative understanding of the growth of science (Dasgupta et al. 1994; Price 1965, 1986). More recently, scholars have analyzed the distribution of citations to understand differences between papers with low or high citations (Iaria et al. 2015; Redner 2004; Thompson et al. 2005; Wang et al. 2013). Part of the literature has also focused on the links between science and innovation by relating patents and scientific articles (Trajtenberg et al. 1997). Using Elsevier's web interface[19], I download the full-text of scholarly articles published between 1970 and 2000 in journals indexed by ScienceDirect[20]. First, I collect the ISSN number of each journal[21], and use the ScienceDirect API to obtain the DOIs of all the articles for each ISSN[22]. Then, I query the full text of articles for the DOIs returned[23]. After a series of cleaning procedures, I obtain a total number of articles of 1,811,301. I detect every document's language and drop non-English articles. Because the translation of

[17]To associate patents to specific countries, I use the country of the assignee. When patents have no assignee but only inventors, I use the country of the inventor.

[18]`http://www.nber.org/patents/`

[19]http://dev.elsevier.com/

[20]I select journals in the following disciplines: chemistry, chemical engineering, engineering, environmental science, materials science, and physics and astronomy.

[21]I do this using Elsevier's website (`https://www.elsevier.com/solutions/sciencedirect/content/journal-title-lists`.

[22]A DOI, or Digital Object Identifier, is a sequence of digits and letters that uniquely identifies an academic article.

[23]Full text data was successfully downloaded for 1,843,684 articles, out of a total of 2,307,345 DOIs initially returned by the API. This implies that Elsevier listed 463,661 DOIs for which the full text was not available. This might be due, for example, to entire journals dropping out of Elsevier Science Direct's collection.

English articles in other languages is often contained in the full text, I drop any sentence containing less than 80% of tokens recognized by a standard English dictionary[24]. More details on the cleaning procedure are provided in Appendix A. For data on affiliations and citation counts, I query the Scopus search API[25]. Additionally, I use the Global Research Identifier Database[26] (GRID) to classify authors' affiliations (e.g., education, company, etc...).

3.2 Tracking Research and Innovation Efforts on Alternatives to CFCs

CFCs are a group of compounds very specific from the point of view of their molecular structure: they contain only carbon, chlorine and fluorine atoms (and typically no more than three carbon atoms). It is precisely thanks to that structure that CFCs have great thermodynamics properties and became broadly used in many different industries. This intricate relationship between molecular structure and industrial properties also implied that there wasn't an infinite number of potential good substitutes. Mainly, those potential substitutes presented similar alkane chains but with hydrogen atoms replacing halogens. For example, one chlorine atom in CFC-12 is substituted with a hydrogen to constitute HCFC-22 or with a methyl group in which case we obtain HCFC-142b (see Figure 1). Here "HCFC" stands for hydro-chlorofluorocarbons. When all chlorine atoms are substituted with hydrogens, the compounds are then known as HFC, or hydro-fluorocarbons. For example, when the two chlorine atoms in CFC-12 are replaced by hydrogens, we get HFC-32.

Strategies for reducing CFCs could include using entirely new designs (like pump-action sprays instead of aerosols) or recycling. But when it came to getting a chemical substitute, scientists knew that the search lied in the realm of HCFCs and HFCs. These molecules had been known for a long time, at least in the lab. The first-ever granted patents related to HCFCs and HFCs typically go back to the 1930s; at the time, chemists were experimenting with halogenation processes and heat transfers[27]. Hence, developing CFC substitutes was not so much about "new-to-the-world" com-

[24] I use SpaCy's English dictionary in Python.

[25] Because of quota limitations, I queried meta-data only for articles mentioning only CFC substitutes, Annex A or B compounds. It is still in progress for articles mentioning HAPs.

[26] `https://www.grid.ac/`

[27] For example, in 1934, a patent is claimed for a "method of producing refrigeration which comprises evaporating in the vicinity of a body to be cooled and subsequently condensing CH2ClF." US Patent 1,968,049. CH2ClF is a.k.a. HCFC-22.

pounds, but instead "new-to-the-industry" compounds. Indeed, the key technological challenges were about making large-scale production cost-efficient, redesigning processes and equipment already installed, and learning about environmental acceptability and human toxicity.

I construct a list of potential substitutes using historical records. After Montreal, manufacturers from the US, Europe, and Japan received authorization from antitrust officials to organize cooperation, at least on the science for which patenting was not possible. They launched two working groups to study the feasibility of various alternatives. The PAFT (Program for Alternative Fluorocarbon Toxicity Testing), created in January 1988, worked on assessing the toxicity of five possible alternatives. The AFEAS (Alternative Fluorocarbon Environmental Acceptability Study), created in December 1988, investigated the atmospheric dynamics of twelve potential CFC substitutes. I use these twelve molecules to form a first group. I also include in this group two other possible CFC substitutes mentioned in Benedick (2009) and Parson (2003)[28]. Table B1 in Appendix B shows the name and additional information about these molecules.

I search through every patent and article to find the documents in which the name of these molecules appear. Being able to search the full text of the documents is an advantage here since relying on abstracts only could lead to many false negatives. Chemical compounds, however, are often given several names; for example, HCFC-22 has 39 other possible names such as chlorodifluoromethane or algeon 22. To capture all the occurrences of a mention of a molecule, I develop an automatic script to collect all possible names for a given molecule through SciFinder, a database of chemical information maintained by the American Chemical Society[29]. I then search through all patents and articles to identify the documents in which any of the names appear[30]. When a document contains the name of only one of the molecules, the document is assigned to that molecule. When it mentions several molecules, it is assigned to each of these molecules[31]. I develop alternative rules as robustness checks. I proceed similarly to identify the patents and articles that mention any of the 171 Hazardous Air Pollutants (HAPs). I explain in the next section how

[28]HFC-245fa and HFC-365mfc are mentioned as possible substitutes in foams.

[29]A full list of all the possible names of CFC substitutes is shown in the supporting online material.

[30]I look for any English name listed in SciFinder but I do not look for chemical symbols. The articles' text is usually the output of optical character recognition, and chemical symbols and formulae are too often rendered with mistakes.

[31]This is what I refer to as the *weak* rule.

these molecules are useful for my methodology[32].

3.3 Topic Proportions

I use topic modeling, a machine learning method for text analysis (Blei 2012; Blei et al. 2006, 2009; Roberts et al. 2014; Roberts et al. 2016), to generate covariates that describe the semantics surrounding molecules. These covariates help describe and measure molecules' chemical and industrial characteristics. The procedure outputs document-level *topic proportions*, that is a variable from 0 to 1 indicating to what extend topic i is present in a particular document. Specifically, I use the algorithm to discover five topics[33]. I then aggregate topic proportions at the molecule level by calculating unweighted and weighted means with weights proportional to the number of times an article mentions a molecule. Figure 2 summarizes these various steps with a simple example of three documents, two molecules, and two topics. Appendix A2 provides a more detailed description of the procedure and the obtained topics.

4 Methodology

4.1 A Sharp Post-1987 Increase

A simple approach to study patterns in patents and articles is to observe the yearly count of documents about CFC substitutes and test whether there is a change of patterns before and after the date of the signature of the agreement. Figure 3 plots the yearly number of articles or patents mentioning the names of any of the 14 CFC substitutes[34]. We note a clear increase after 1987, the year Montreal was signed. The hypothesis is that the signature of Montreal acted as a strong signal to the business community that prices were going to change and modified expectations regarding where future profits lay, i.e., in CFC substitutes[35]. We shall add that the response to an

[32]The full list of HAP molecules that I consider is available in the supplemental online material.

[33]I run robustness checks with ten topics. Increasing the number of topics will output finer grained topics. However, implementing the SCM becomes then more computationally challenging.

[34]Figure B2 in Appendix B displays similar trends for each of the 14 molecules.

[35]I shall highlight here that the protocol was implemented in the USA through the 1990 Clean Air Act. Specifically, the reduction targets agreed at the international level were imposed to each firm known to be either a producer or an importer of CFCs. The reductions were calculated from each firm's baseline level of production

international agreement should be greater if many countries participate in that agreement(Dekker et al. 2012). A global market means more profit-making opportunities for firms about to incur the sunk costs of research and development. For researchers publishing in peer-reviewed journals, the Montreal Protocol likely acted as a strong signal that a technological transition was underway inducing researchers to redirect their work towards CFC substitutes. Additionally, the choice of which scientific research to conduct can also be heavily influenced by organizations funding research, such as the National Science Foundation. Likely, research grants were directed on ozone depletion and alternatives to CFCs. These mechanisms could potentially explain why we observe a strong increase in patents and articles on CFC substitutes after 1987[36].

4.2 A Description of Patents and Articles mentioning CFC substitutes

Table 1 illustrates what kind of patents mentioned CFC substitutes: the table lists the ten most common patent codes. Here, the codes refer to the International Patent Classification. Unsurprisingly, we see that most codes belong to the C class, the class for Chemistry and Metallurgy. We find many in the subclasses "C07" and "C08" which refer to the purification, separation or stabilization of organic compounds possibly containing carbon and halogens with or without hydrogen. For example, C07C 19/00 corresponds to "Acyclic saturated compounds containing halogen atoms". Additionally, Figure B1 in Appendix illustrates that the patent codes most frequent before 1987 tend to also be the most frequent after 1987. At the same time, some codes with low frequency before 1987 become important after 1987[37].

Table B2 in Appendix B displays summary statistics about countries and affiliations of patent assignees and authors of articles. The typical patent is granted to a for-profit organization in the United States. Indeed, more than 96% of patents are granted to for-profit organizations while

where the baseline year was 1986. The Clean Air Act also specified that, starting in 1992, the EPA would issue lists of acceptable and unacceptable substitutes (a.k.a the SNAP initiative); from then on, it was only allowed to substitute CFCs with approved substances. The first list of acceptable and unacceptable substitutes was published in 1994.

[36]The analysis of NSF grants is in progress.

[37]C08G: Macromolecular compounds obtained otherwise than by reactions only involving carbon-to-carbon unsaturated bonds, C10M: Lubricating compositions; Use of chemical substances either alone or as lubricating ingredients in a lubricating composition, C23G: Cleaning or de-greasing of metallic material by chemical methods other than electrolysis, C11D: Detergent compositions; Use of single substances as detergents; Soap or soap-making; Resin soaps; Recovery of glycerol.

the rest is shared among organizations coming from the educational and governmental sector. European assignees represent only around 20 to 30% of patents; Japanese around 10 to 20%. Figure B3 in Appendix illustrates that the increase in the number of patents mentioning CFCs applies to all countries; we note a particularly strong increase for patents with assignees located in Japan and the UK.

Table 8 displays the titles of the most cited articles mentioning CFC substitutes. Only articles with three molecule occurrences in the text were kept in the sample. We note that these articles, as expected, seem to focus on chemical and physical characteristics of CFC substitutes ("boiling", "evaporation", "pressure" etc...).

4.3 Was Innovation a Key Driver to the Agreement?

An interesting question is whether R&D activities before 1987 eventually led to the success of the negotiations at Montreal. On Figure 3, the trend in patenting and publishing before 1987 looks not just astonishingly flat but also very small in terms of actual number of patents and articles published each year. This is indicative that, in fact, little was going on before Montreal on the science and innovation on CFC substitutes. The ozone crisis literature often mentions the existence of domestic regulations before 1987 as potential pre-Montreal drivers of innovation. For example, in August 1977, the U.S. Congress amended the Clean Air Act with the Stratospheric Ozone Protection amendment writing into law a CFC ban on aerosols by 1978[38]. In fact, even before any domestic regulation, some manufacturers unilaterally decided to remove CFCs from their spray products because they worried about their public image. Hence, consumer pressure possible acted as an incentive for firms to innovate. However, these pre-Montreal domestic regulations and unilateral actions on behalf of manufacturers only targeted aerosols, one very specific industrial application of CFCs for which substitutes[39] could easily and cheaply be implemented[40]. These product changes unlikely required a significant research effort. The low levels in patent and article count between 1970 and 1987 on Figure 3 indicates that neither consumer pressure nor

[38]Similarly, in 1978, Canada, Switzerland and Scandinavian countries all banned CFC aerosols. On continental Europe, Germany called for a European Community-wide ban without success.

[39]Physical substitutes included roll-on devices; chemical substitutes included alkanes.

[40]In 1980, the EPA proposed to freeze other uses beyond aerosols but U.S. industry blocked the initiative.

aerosol regulations seemed to have stimulated science and innovation on the 14 CFC substitutes I consider[41].

The flat trend in patenting and publishing before 1987 on Figure 3, additionally, indicates little anticipation of the negotiations' success. This is also supportive of the negotiations being little influenced by the research output on CFC substitutes. Indeed, if R&D output were key drivers to the negotiations' success, firms should have been able to anticipate the negotiations' outcome. Firms have strong incentives to be forward-looking because anticipating can confer a first mover advantage. By undertaking early-on research activities, they can develop cleaner technologies before competitors and build a strategic advantage when regulations are passed[42]. As a result, if the signature of Montreal had been anticipated, we would observe a gradual increase in patent and article counts starting before or at least close to 1987. It is difficult to pinpoint precisely the optimal patenting timing. On the one hand, firms have an incentive to delay patenting right until production and commercialization begin because patents expire after ten years. This phenomenon is particularly salient for technologies with little risk of a competitor developing a comparable product. Patenting renders the output of a firm's R&D public knowledge, allowing potential competitors to effectively learn from it and come up with even better technologies. On the other hand, when competitors work on closely related projects, delaying patenting sharply increases the risk that competition patents first. This mechanism was likely salient in the case of CFC substitutes. Hence, overall, we should not expect CFC manufacturers to delay much of their patenting activity. Instead, it seems firms believed the likelihood of regulations to be low and therefore had little incentive to invest in R&D.

This is, in fact, consistent with Benedick's accounts of the events (Benedick 2009): when the issue took prominence in the media and regulators' minds in the late 1970s, firms initiated some R&D projects regarding CFC substitutes. However, those projects were canceled by the early 1980s when the probability of regulation rapidly converged to zero. At that time, uncertainties in the science of atmospheric ozone seemed irreducible, and the year 1981 saw the election of a strongly

[41]These 14 CFC substitutes were targeting foams, refrigeration and solvent applications of CFC.

[42]Firms with such competitive advantage could even decide to lobby in favor of environmental regulations for that reason (Puller 2006).

anti-regulatory American administration. In Europe, many governments persisted in refusing to harm their domestic manufacturers with any regulation. Benedick emphasizes the complexity of the negotiations and the great uncertainty, until the last minute, of the negotiations' outcome. He further argues that some exceptional turns of events unlocked the situation. Unexpectedly, Reagan overruled his own administration and approved the agreement: the U.S. President had a skin cancer removed twice in the past, and Benedick hints that Reagan's life experiences weighed heavily on his decision. On the European side, the biggest opponent to the regulation of CFCs, the U.K., left the European Community Presidency, leaving Germany, Denmark, and Belgium, firm proponents, as the head negotiators. This account of the negotiations' success does indeed indicate that the agreement largely occurred independently from the state of R&D activities on CFC substitutes.

Consequently, it looks unlikely that science and innovation on CFC substitutes were strong drivers of the diplomatic efforts to regulate CFCs in Montreal. At this stage, I can not rule out that firms might have kept their CFC substitutes secret and discretely lobbied for the success of the agreement. I will investigate this possibility in Section 6.

4.4 First Differences

To quantitatively investigate the temporal patterns and detect a change happening in 1987 onwards, I implement the following econometric specifications: a first difference specification with a mean shift (Equation 1) and a first difference specification with a trend-break (Equation 2). $LogCount_{m,t}$ is the log number of documents in year t about molecule m; $\lambda_{post1987}$ is a dummy variable that equals one when $t > 1987$; λ_m are molecule fixed effects; $Years$ is a continuous variable indicating the number of years relative to 1987. Here, I suggest to use counts in log, instead of level, as the outcome variable since it will provide a better linear fit over time. Indeed standard models suggest we can think of scientific production as exponentially growing over time. The sample here consists of 14 different CFC substitutes of which I track the number of patents

and articles throughout the years.

$$LogCount_{mt} = \alpha + \beta_0 * \lambda_{post1987} + \lambda_m + \epsilon_{mt} \tag{1}$$

$$LogCount_{mt} = \alpha + \beta_1 * Years * \lambda_{post1987} + \beta_2 * Years + \lambda_m + \epsilon_{mt} \tag{2}$$

The main hypothesis is that β_0 and β_1 are both positive for CFC substitutes, implying a significant increase in research and patenting activities relating to CFC substitutes after 1987 once Montreal passed. Table 2 displays the regression tables for the simple first time differences. Model 1 confirms that there is a significant and positive mean shift after 1987 in the number of patents and articles mentioning CFC substitutes. The coefficients corresponds to about 630% more patents and 190% more articles on CFC substitutes after 1987 (compared to before). Model 2 shows that the change can also be modeled as a trend break. The coefficient for "Years" indicates that there is a small positive underlying trend for both patents and articles.

Because this is only a simple temporal difference, such an increase could also be due to other underlying trends not specific to ozone negotiations. For example, it might be possible that some other reforms or the economic context fostered more academic and industrial research in the 1990s. Hence, we need to find a group of molecules that could serve as a control group; the challenge consists in finding molecules that are very similar to the treated molecules, while, at the same time, remaining different enough to ensure that they are not affected by the treatment. Specifically, a good control group should contain molecules which undergo similar influences as CFC substitutes apart from the one of the Montreal Protocol. One way of choosing such molecules is such that they present similar pretreatment trend in the outcome variable *LogCount* but also such that they are as close as possible to the treated molecules chemically, physically and regarding industrial applications. Such molecules can potentially be found in the pool of HAPs.

5 Difference-in-Differences and Synthetic Control Method

5.1 HAPs as a Comparison Group

HAPs have no connection to ozone but they are often related to industrial activities. They became monitored under the Clean Air Act due to human health concerns including cancer, asthma, birth defects, reproductive effects, and neurodevelopmental effects, as well as adverse ecological impacts. Examples include benzene, chromium or formaldehyde[43]. Figure 4 illustrates why HAPs are a good choice as control molecules: overall patents about CFC substitutes and HAPs fall into similar top-level codes. Additionally, Figure B4 shows that they also display similar second-level patent codes. Table B2 in Appendix B displays summary statistics about countries and affiliations of patent assignees and authors of articles[44]. The two groups have similar profiles with almost all patents being granted to a for-profit organization, and about half being domiciliated in the United States.

In 1990, an amendment to Clean Air Act required the EPA to promulgate regulations establishing emission standards for each listed category of major sources and area sources of HAPs. The EPA published the initial list of "source categories" in 1992, that is the list of industries and production processes targeted by the regulations. In 1993, the EPA published an initial promulgation schedule specifying by which year[45] sectors were expected to comply with the emission standards [46] for each category or subcategory of major sources and area sources of HAPs. Importantly, the standards required the maximum degree of emission reduction that the EPA determines to be achievable by each particular source category. This is know as MACT, or maximum achievable control technology[47]. In other words, the rest of the industry adopts what already exists, and the incentives to innovate are limited. Typically, firms install end-of-pipe pollution control such as destruction (through thermal oxidation, catalytic oxidation, flaring) which is useful when the

[43]The full list of the molecules included in the different treatment groups is displayed in the supporting online material.

[44]Data collection for HAPs in articles is still undergoing due to quota limitation on the Elsevier API.

[45]Most sectors were asked to comply by 1997 or 2000. A few only by 1994.

[46]a.k.a. the National Emission Standards for Hazardous Air Pollutant (NESHAP)

[47]Less stringent standards, known as generally available control technology (GACT) standards, are allowed at the Administrator's discretion for area sources.

HAP is a waste gas, or through recuperation (using diverse methods such as adsorbers, absorption/scrubbing, concentrators, condensation, biofiltration, membrane technology etc...). Additionally, much emissions reduction came from limiting fugitive emissions from storage tanks and better management practices such as leak detection and repair (Moretti 2002).

On Figure 5, I plot the yearly mean counts of documents mentioning CFC substitutes and HAPs to check similarity in the pre-trends. When using the whole sample of HAPS (that is 171 HAPs molecules), pre-trends look similar. However, a close observation reveals that, in patents, pre-trends for CFC substitutes are slightly up, while the one for HAPs goes slightly down. In articles, HAPs have a clear upward trend before 1987 while CFC substitutes seem somewhat flatter. Similarity in pre-trends can be improved by selecting a subset of HAPs with pre-trend closest, in terms of log count, to the average CFC substitutes. Specifically, I construct the DiD control group such that it contains the 42 HAPs whose pre-trend is closest to the average trend of CFC substitutes[48].

5.2 Difference-in-Differences

I estimate the DiD model with a mean shift specification (Equation 3) and a trend-break specification (Equation 4). $LogCount_{m,g,t}$ is the log number of documents with molecule m belonging to molecule group g, in year t; $Post_t$ equals one when $t > 1987$; D_m equals one if the molecule belongs to the treated group; λ_m are molecule fixed effects; λ_t are year fixed effects; $\boldsymbol{X}_{mt}$ is a vector of covariates; Y is a continuous variable indicating the number of years relative to 1987. β_0 identifies the DiD estimate.

$$LogCount_{mt} = \alpha + \beta_0 \cdot D_m \cdot Post_t + \lambda_t + \lambda_m + \boldsymbol{\gamma_t} \cdot \boldsymbol{X}_{mt} + \epsilon_{mt} \tag{3}$$

$$LogCount_{mt} = \alpha + \beta_1 \cdot Y \cdot Post_t \cdot D_m + \beta_2 \cdot Y \cdot Post_t + \beta_3 \cdot Y + \lambda_t + \lambda_m + \boldsymbol{\gamma_t} \cdot \boldsymbol{X}_{mt} + \epsilon_{mt} \tag{4}$$

[48]First, I calculate the mean pre-trend slope for CFC substitutes between 1976 and 1985 (it is equal to âĹŠ0.018). I then calculate the pre-trend slope for each HAP between 1976 and 1985. I rank HAPs according to how close their slope is to the mean slope for CFC substitutes. Finally, I select the 42 HAPs with most similar slope. I construct the control group such that it is three times larger than the treated group. There are 14 molecules in the treated group, hence I use 42 units in the control group (3×14).

The primary hypothesis is that β_0 and β_1 are positive. Significant coefficients would imply that the research and development activities underwent important changes after 1987 relative to the counterfactual. If there is no significance, this might suggest that the research effort was redirected towards CFC substitutes already before the signature of the treaty.

Table 3 displays the DiD results. Model 1 corresponds to the main differences-in-differences specification. It includes year and molecule fixed effects. The binary variable "Post 1987 x Substitutes" equals 1 for observations belonging to the group CFC substitutes and after 1987. For patents, the coefficient is smaller than the coefficient in the simple difference, but remains significant and large, corresponding to more than a 500% increase. This estimate corresponds to an additional 28 documents per year for the average substitute. Since there are 14 CFC substitutes in my sample, this implies 390 additional patents a year for CFC substitutes in aggregate. For articles, the coefficient is smaller than the coefficient in the simple difference, but remains significant and large, corresponding to more than a 95% increase. This corresponds to an additional three documents a year for the average substitute; hence, aggregating the 14 CFC substitutes, this corresponds to about 47 additional articles per year. Model 2 presents a trend-break specification. It shows that the log number of patents mentioning CFC substitutes increases with the years after 1987 by 0.22 more than the control group. Similarly, the log number of articles mentioning CFC substitutes increases with the years after 1987 by 0.10 more than the control group.

Figure 6 display the DiD coefficients plots. We note that, in patents, the treatment effect is small, yet statistically significant, as early as 1988. For articles, the treatment effect is first statistically significant in 1990. Two mechanisms can account for a delay between the moment firms and researchers decide to redirect their efforts towards CFC substitutes and the granting of a patent or publication of an article. First is the lag between application and granting of a patent; and similarly the lag between submission and publication of an article. The current analysis already reflects such lag for patents (since the date used is the patent application date)[49].

[49]The average delay between application and granting for patents on CFC substitutes is 22 months (with a standard deviation of 12 months). It is very similar for patents on HAPs (about 23 months, with standard deviation of 12 months). Figure B5 in the appendix graphically shows the differences between application and granting date. Overall, plotting the number of patents based on granting dates simply shifts the entire curve two years forward.

However, the time-series of articles reflect their year of publication.

Second, the time required to obtain any technology worth patentable can broadly vary. It is difficult to assert how long it takes firms to develop new technologies in response to a change in incentives. We can expect such delay to vary from technology to technology even within the same technological sector. In the context of energy patenting, Popp (2002) estimates that the mean lag occurs in 3.71 years and the median lag in 4.86 years. This implies that over one-half of the full effect of an energy price increase on patenting is experienced after just 5 years. These estimates are somewhat consistent with the shape of the yearly treatment effects obtained here.

One possibility is that firms developed technologies prior to 1987 keeping them secret. Once Montreal is agreed, it becomes worthwhile to patent as firms know they will eventually commercialize them. It is possible, therefore, that some patents granted soon after 1987 results from R&D effort incurred prior to Montreal. I will further investigate this question in section 6.

I run additional DiD specifications controlling for lags of log count and for topic proportions. Table 4 compares the mean numbers of documents and topic proportions for patents and articles across CFC substitutes and HAPs. The two groups have very different average counts, and mean topic proportions are also statistically different across the two groups. Table 5 indicates that the treatment effects remain robust to those control variables. The magnitude of the treatment effects, however, is reduced.

5.3 SCM and DiD

DiD strategies are designed to estimate average effects over a population from which we sample a large enough number of units exposed and units non-exposed to treatment. Considering the overall population of potential CFC substitutes, I have sampled 14 of them; however, those 14 observations are not independent because several molecules are often mentioned in the same documents. Additionally, the reported standard errors in the DiD regressions reflect uncertainties about the aggregate data. This is problematic because the greatest uncertainty lies in the choice of the control group, and not in the aggregate quantities. In fact, here, the aggregate quantity can be thought of as observed: adding up the observed counts of the 14 CFC substitutes, and

considering them as one single treated unit. Figure 7 plots the number of patents (in log) mentioning CFC substitutes. The thick line called "Substitutes (aggregated)" corresponds to the number of patents mentioning any of the 14 CFC substitutes. I implement the SCM on this aggregated substitute and I am interested in examining whether the aggregate count of these 14 substitutes has gone up compared to a control group. SCM was specifically developed to evaluate the effects of large aggregate interventions when the treatment affects an aggregate quantity (Abadie et al. 2003, 2010, 2015; Athey et al. 2016). Many interventions are in fact implemented at an aggregate level and have an impact on a small number of large entities, such as cities, school districts, or states. I enlarge the application of SCM to a new kind of aggregate entity: field of scientific and engineering inquiry.

The magnitude of the treatment effects estimated with the DiD strategy inherently relies on the choice of the control group. It is possible to improve those estimates by choosing the molecules that are included in the control group more precisely. Figure 8 in the appendix shows patent counts in log for each HAP and for the aggregated CFC substitutes. The graph illustrates the high heterogeneity within the group of HAPs, and in particular that many HAPs have log counts much higher than the aggregated CFC substitutes. Some of the HAPs might also be very different chemically, physically and from an industrial point of view. An improvement on using all HAPs in the control group is therefore to use only the HAPs whose yearly number of patents and articles are driven by structural processes most similar to those driving those of CFC substitutes. The selection of comparison units is crucial in such study: if too different from the treated unit, any deviation in the outcome after the treatment can be attributed to initial differences and the resulting estimate would be biased. The SCM offers a data-driven way to construct comparison units using only the HAPs most similar to CFC substitutes. Figure 9 illustrates why topic proportions are useful in this case. We see that some HAPs have values of topic proportions that stand out as outliers, indicating that those HAPs present a semantic context that is likely very different from the one of CFC substitutes. Using topic proportions together with the SCM ensures that such HAPs are not used in constructing a synthetic control.

The key idea of SCM consists in using a weighted average of a set of control units with the

weights chosen so that the weighted average is similar to the treated unit regarding covariates and outcome in the pretreatment periods[50]. The advantages of SCM relies on the opportunity to create a synthetic unit that shares as much as possible the characteristics of the treated unit. In my case, I would hope to construct a synthetic unit that not only reproduces the path of counts in pretreatment periods, but that also resembles the treated molecules regarding chemical, physical and possibly industrial characteristics. To this aim, I attempt to proxy such characteristics with topic proportions derived from topic modeling of the documents' text. I explain this further in details in the next section.

5.4 Implementing SCM

I implement the SCM using log count as the main outcome variable considering the 14 CFC substitutes as one treated molecule. The outcome variable is therefore the log number of patents or articles that mention any of the 14 molecules. The synthetic control is constructed by fitting the values of log counts in the pretreatment periods and the topic proportions[51]. In all the SCM specifications, the treatment year is the first year in which the treatment becomes active: this is defined as 1988 since Montreal was agreed in 1987. To be conservative, I use data from 1970 until 1985 only to fit the synthetic control[52].

It is critical to ensure that the synthetic control closely matches the treated unit in the pre-treatment periods. If that was not the case, the synthetic control unlikely provides a good proxy of a counterfactual since it is not even a good proxy of the treated unit before treatment. Following Abadie et al. (2010), I examine the Root Mean Square Prediction Error for periods before treatment (pre-RMSPE)[53] to verify whether the discrepancies between the synthetic control and the treated unit are large and thus whether the SCM is appropriately implemented.

[50]For example, suppose we had 3 HAPs as control units with weights μ_a for asbestos, μ_b for benzene, and μ_c for catechol. Then the weights are chosen such that $\mu_a \times Y_{at} + \mu_b \times Y_{bt} + \mu_c \times Y_{ct}$ is close to Y_{St} (where S stands for substitutes) for periods t before the treatment takes place. Here, Y is the log count of articles mentioning the molecule, but additional covariates can be used.

[51]I run different specifications: with weighted means and with unweighted means.

[52]Topic proportions are averaged over the entire pre-1985 period, while log count is not.

[53]The pre-RMSPE measures lack of fit between the path of the outcome variable for any particular unit and its synthetic counterpart: the pre-RMSPE of unit 1 is defined as $(\frac{1}{T_0}\sum_{t=1}^{T_0}(Y_{1t} - \sum_{j=2}^{J+1} w_j^* Y_{jt}))^{1/2}$ where T_0 is the number of pretreatment periods. A post-RMSPE can be similarly defined for periods going from $T_0 + 1$ to the end of time-series available.

I first use the 171 HAPs as units in the donor pool. However, as explained by Abadie et al. (2015), reducing the size of the donor pool can limit the risk of over-fitting as well as the risk of interpolation biases. Following their advice, I use a smaller donor pool containing only the HAPs that are close to the treated unit in the space of covariates and outcome. I choose the twenty HAPs with lowest pre-RMPSE, that is the twenty HAPs that are closest to the treated unit in terms of topic proportions and count. In what follows, I call this group of HAPs the "smaller" donor pool. I refer to the "whole sample" of HAPs when the donor pool includes the 171 HAPs. Finally, I also check that there is no risk of extrapolation[54].

For inference, I follow the method suggested by Abadie et al. (2010) and Abadie et al. (2015). The exercise consists in applying the SCM procedure to every potential control in my sample. This allows me to assess whether the effect estimated for the unit treated is large relative to the effect estimated for a molecule chosen at random. This is akin to implementing placebo tests wherein each unit in the control group is assumed to have received the treatment at the year 1987. A synthetic control is then constructed for each placebo, and we observe what would have been the hypothetical treatment effect for this "falsely" treated unit. This creates a distribution of placebo effects, and we can evaluate the effect for the "true" treated unit vis-a-vis to where it falls in this distribution. A p-value is calculated as the fraction of placebo effects that are greater than or equal to the effect estimated for the "true" treated unit.

As suggested by Abadie et al. (2010), it is useful to compute the ratios of post-RMSPE over pre-RMSPE and examine where in the distribution of those ratios, the treated unit lies. For example, if the treated unit is second largest ratio among a donor pool of 50 units, then the p-value can be computed as $\frac{2}{50} = 0.04$, and the treatment effect would be significant at the 5% level[55]. The p-value can be interpreted as the probability of obtaining an estimate at least as large

[54]When choosing weights for the donor units to create the synthetic control, the SCM algorithm imposes that the weights sum to 1 and that they be nonnegative. These constraints avoid any risk of extrapolation. However, when the treated unit presents values for covariates that are either the smallest or the largest in the distribution of the donors, it becomes difficult to approximate it. To verify that the donor pool remains adequate, Table B3 presents summary statistics for CFC substitutes and the small pool of HAPs in the case of counts derived from the weak rule. It is reassuring to see that the range of values displayed by the HAPs always contains the value for CFC substitutes. Hence, here, the constraints that weights must sum to 1 and be non-negative does not seem to be an issue.

[55]The treatment effect's p-value for the treated unit is therefore defined as: $p_1 = \frac{\sum_{j=2}^{J+1} 1\{ratio_1 \geq ratio_j\}}{J}$, where

as the one obtained for the "true" treated unit. Hence the inference is mostly limited to assessing whether the treated effect is large compared to the distribution of the placebos. To illustrate, how p-values are calculated, Figure 10 displays the distribution of post-RMPSE over pre-RMPSE for the case of log count, weighted means of topic proportions and the whole sample of HAPs, for the corpus of patents. The figure shows that the ratio for CFC substitutes is greater than all of the 168 other units. Hence the p-value, in this case, is 1/168.

5.5 SCM Results

Before examining results, I want to illustrate the benefits of the SCM. Table B5 compares the mean value over the years 1970 to 1985 of log count and topic proportions for CFC substitutes, the observed treated unit ("real S" in the table), for the comparison unit constructed through the SCM ("Synthetic S") and for the average of HAPs[56]. We see that the synthetic control matches the "real" CFC substitutes group better than the average of HAPs in terms of log count. This is the core idea motivating the use of the SCM[57]. Table B4 illustrates that topic proportions contribute around 15% in constructing the synthetic control[58].

Table 6a summarizes the performance of the main SCM implementations for CFC substitutes in the corpus of patents. I ranked the table according to the magnitude of the pre-RMPSE with smaller pre-RMPSE at the top of the table. The lower the pre-RMPSE, the better the fit between the synthetic control and the treated unit over pretreatment years, and therefore the more credible it is that the synthetic control appropriately proxies the counterfactual. Furthermore, recall that procedures using the small pool are more trustworthy because they limit the risk of interpolation biases and overfitting. Hence the preferred specifications here uses the small pool and the weak rule and weighted means of topic proportions. The p-values smaller than 0.01 indicate significance at the 99% level. Table 6a also reports the year in which the treatment begins to be significant

$ratio_j = \frac{post-RMPSE_j}{pre-RMSPE_j}$ and subscript 1 refers to the treated unit.

[56]The HAPs used in calculating the average are only those from the small pool. The synthetic control, here, was constructed based on similarity with the variables "Log count" and the weighted means of the five topic proportions.

[57]The topic proportions are very similar, but this could be expected since the pool of HAPs used is the small one, and the means of topic proportions within that pool are very concentrated.

[58]In the Stata *synth* package, these weights are determined according to the amount of predictive power that each variable has over the outcome. Hence, in the case of patent counts with weighted means of topic proportions, the outcome variable, log count, is the variable assigned greatest weights.

at the 10% level. This is determined by calculating p-values for each year separately. We see here that the treatment effect becomes significant in most cases only starting in 1990 or 1992. The ATE for the preferred specification equals 0.89; this corresponds to a 144% increase in patents compared to the synthetic control; this gives a treatment effect close to 117 patents per year.

Similarly, Table 6b provides performance summary for the main SCM procedures for CFC substitutes in the corpus of articles. Here, the preferred procedure to report the ATE uses the small pool of HAPs and weighted means of topic proportions. The ATE for this preferred specification equals 1.06 and is significant at the 99% level. This estimate translates to a 189% increase, and corresponds to about 43 additional articles per year.

Figure 11 graphically displays the results of the SCM for CFC substitutes. The graphs correspond to SCM implementation for the preferred specification, that is when the SCM is implemented with log counts, using the small pool of HAPs and weighted means of topic proportions. The graphs on the left-hand side represent the raw effect, that is the observed time series of the treated group along with the time series of the constructed control. On the right-hand sides are shown the placebo tests, the non-parametric tests to evaluate the significance of the results; the black lines show the effect on the treated group relative to the control group, while each gray line is a placebo test performed on a unit drawn from the donor pool.

The figure illustrates what we concluded from Table 6b and 6a: the treatment effect on CFC substitutes appears significant for both patents and articles. We note that the black line rises above most other lines mostly as from 1990. This indicates that, similarly as in the DiD, the treatment effect is statistically significant only after 1990.

5.6 Verifying SCM Assumptions

Anticipation. An important assumption supporting the SCM is that the intervention does not affect the outcome before the implementation period. In reality, anticipation effects often violated this assumption; part of the treatment effect would become embedded in the control, and the SCM would lead to understating the treatment effect. A workaround consists in redefining the treatment year as the first period in which the outcome may react to the intervention. I have

already discussed that anticipation seems unlikely. However, I nonetheless replicate the SCM using 1985 as the beginning of the treatment since it was in March of that year that the Vienna Convention was adopted. The meeting in Vienna can be considered as the start of the ozone layer's diplomatic life. Figure 12 displays the SCM graphs for the preferred specification: it uses the small pool and weighted means of topic proportions for both patents and articles. Here, the earliest possible take-off would be in 1983 since I fit the synthetic control using data up to 1982. We observe that there does seem to be no takeoff before 1987 and results are very similar to fitting the synthetic control up to 1986.

Interferences A second assumption supporting the SCM requires that there be no interferences between units, meaning that HAP molecules should not be affected by the Montreal Protocol. This is unlikely to be the case since HAPs have been under the regulatory radar for very different reasons than ozone depletion. However we may be worried that the redirection of research efforts towards CFC substitutes crowded out efforts towards the control molecules. This is also unlikely since HAPs are used in different types of industrial activities. However I still proceed to a careful examination of the firms patenting both on CFC substitutes and the HAPs contributing to the synthetic control. Table 7a and 7b provide a description of the HAPs entering the synthetic control for patents and articles, respectively[59]. We note that many of the industrial applications are not directly related to those of CFC substitutes which indicate a crowding out is unlikely. I investigate to what extend the assignees of patents on CFC substitutes and on those HAPs are similar[60]. I find that about 60% of patents mentioning CFC substitutes after 1987 are issued to assignees that never patented on any of the synthetic control HAPs[61]. Examples of such assignees are firms like 3M, Allied Chemical, BASF, Dow Chemical and Procter & Gamble.

[59] A longer description is available in Table B10.

[60] Unfortunately, assignee names in patent records are not standardized and the same firm can appear under different variations of the same name. I therefore use a fuzzy matching algorithm in order to match assignee names.

[61] There are a total of 535 different assignees patenting on CFC substitutes after 1987, and 125 of those assignees about 25% also detain patents on one of the synthetic control HAPs.

5.7 Robustness Checks

Greater number of topics I increase the number of topics generated by the LDA topic model from five to ten to allow for finer grained topics[62]. I, therefore, use ten different topic proportions as covariates in the SCM procedure. Results are displayed in Table B8 in the appendix. I find similar treatment effects.

Are the HAPs picked up by the SCM the same accross various specifications? Table B9 shows which HAPs was selected to construct the synthetic control under the different specifications. First, we note that the HAPs selected for patents are usually different from those selected for articles. This reflects the fact that the contents of patents and articles indeed differ. Second, the same HAPs are picked up by the SCM accross different specifications. For example, 3,3-Dimethoxybenzidine is selected accross the four specifications in patents. Similarly, the procedure with patents typical picks up one of the cresol compounds[63].

Molecule Frequency Here, I check whether the results are robust to dropping patents and articles that mention molecules only once or twice or three times etc... Figure 13 illustrates that it would change little of the analysis. Indeed the trend of the average HAPs remain very similar; only the levels decrease as we increase the threshold of occurrence. If anything, using on patents with only a greater number of occurrences seems to exacerbate the differential between the pre and post trends for CFC substitutes. Figure 14 shows that the SCM results are robust to using patent counts weighted by molecule occurrences.

Counts in Level I replicate the SCM procedures using counts, in level and not in log, as the outcome variable. Table B6a and Table B6b display the results for patents and articles, respectively. I find an ATE that equals about 111 additional patents every year from 1988 to 2000, significant at the 99% level. This estimate is very close to the treated effect obtained using logged counts, which was about 100 patents per year. For articles, the ATE for the preferred

[62]Increasing the number of topics is likely to improve the coherence of each topic. However, implementing the SCM with a large number of topics would be difficult computationnaly.

[63]Table B10 provides a description for each of those HAPs.

specification equals about 44 additional articles every year from 1988 to 2000, significant at the 99% level. Like in the case of patents, this estimate is very close to the treated effect obtained using logged counts, which was about 40 articles per year.

Other Assignment Rules I consider a different rule for assigning document to molecules to test the robustness of my main results. Under the basic rule, which i call the *weak* rule, a document was assigned to group X if a molecule of group X is mentioned in the document, regardless of whether molecules from other groups are also mentioned. Now, under the new rule, which I call the *intermediate* rule, a document is assigned to group X if the molecule with the greater number of mentions is from group X. Figure B6 in the appendix shows the number of patents and articles every year for each molecule group according to the three assignment rules. We note that the weak rule displays a greater number of documents for CFC subtitutes; for HAPs, the weak and intermediate rules overlap almost completely. Table in the appendix B7 displays the performance results of the SCM procedures for patents and articles, respectively. We see results are very similar to the main specifications. The estimated treatment effects are somewhat larger than those with the weak rule. The pre-RMSPE values for patents however indicate that the synthetic control using the weak rule provided a better fit.

Other Robustness Checks Finally, since I implemented a DiD design using the particular subset of HAPs, I also implement the SCM using that subset as the donor pool. I find results that are comparable to the main specification. Details are reported in the supplemental online material. For patents, the treatment effect found approximates 0.9, while for articles it is lower than 0.80. We note that the values of pre-RMSPE are higher than the ones for similar procedures using the small pool. This implies that the procedure to select the subset based on similar slope excluded some molecules which ended up being useful contributions to the synthetic control.

6 Further Results

6.1 Influential Patents and Articles

In this section, I show that the most influential patents and articles were published after Montreal. Indeed, a possible alternative explanation is that despite most patents and articles being published after Montreal, the most influential ones happened before the signature of the treaty. Figure 15 indicates that, on the contrary, highly cited patents and articles concentrate after 1987[64]. I also implement an SCM using patent counts weighted by the number of citations and find strong treatment effects (see Figure 16). Similarly, Figure 14 shows results are robust to using patent counts weighted by both citations and molecule frequency.

6.2 Secret CFC Substitutes?

Is it possible that firms initiated the transition to CFC substitutes before the Montreal Protocol, without patenting but instead keeping their technologies as trade secrets? Indeed, some firms announced at the end of the 1970s that they started R&D into CFC substitutes. Although the same firms soon after announced they terminated those R&D programs, it has been suggested that they might have developed key technologies that they kept secret. Figure 17 indicates that the post-Montreal burst of innovations on CFC substitutes is not driven by a few firms that would have been historically patenting on CFC substitutes since the 1970s. This finding cast the first doubt on the secret substitutes hypothesis. Specifically, Figure 17a shows that after Montreal there are many more firms with patents mentioning CFC substitutes and HAPs. It indicates the likely presence of new entrants in the post-1987 period. Figure 17b confirms this phenomenon by plotting the yearly number of assignees that are "new", meaning it is the first time they appear in the data with a patent mentioning CFC substitutes and HAPs. The figure shows that, after 1987, many firms with no prior experience on CFC substitutes begin patenting.

Additionally, if secret CFC substitutes existed, we would expect a one-time increase in patent

[64]Graphs 15a and 15c include patents and articles, respectively, that mention at least one occurrence of a molecule. To test the robustness of this findings, I plot similar graphs but for patents and articles that mention at least 3 occurrences in Figure 15b and 15d.

counts in the immediate aftermaths of Montreal. Figure 18 plots the number of patents mentioning CFC substitutes by month in the two years that followed Montreal. We see in the first graph that there is no patenting peak. Furthermore, if the extent of R&D efforts provided before Montreal was the key driver to the post-Montreal increase in patenting, we should observe major differences in the patenting trends of old and new entrants. On the second graph, I present trends for assignees that never obtained any patent mentioning CFC substitutes before 1987 and those who did. Although a gap seems to build up over time, trends look mostly similar.

6.3 Zooming in on Key Manufacturers

Figure 19 illustrates the possible key role of a few manufacturers: the scatter plot shows, for each firm in the sample, the number of patents between 1975 and 1986 on the x-axis and the number of patents in the two years that followed Montreal on the y-axis. We see that a positive trend is mostly driven by three firms: DuPont, Allied, and Dow. Excluding those, there are no clear correlations between patenting before 1987 and patenting in the immediate aftermaths of Montreal. This plot, however, motivates a more detailed investigation in the behavior of DuPont and Dow. Figure 20a shows that most patents granted to DuPont and Allied were applied for after 1989. Figure 20b shows that there is no sudden peak patenting right after Montreal. Instead, we observe a gradual ramping up of patenting activity. Figure 20c illustrates that the patents granted to DuPont and Allied which received the greatest number of citations mostly originate from 1989 to 1991. Figure 20d indicates, however, that, in the weeks that followed Montreal, both DuPont and Allied applied for patents that would go on receiving a high number of citations. This seems to indicate that DuPont and Dow likely had a first mover advantage on some technologies. However, the magnitude of the ramping up in patenting activity that follows from 1990 onwards allows concluding that most of the innovative activity started after Montreal[65].

[65]Another way of examining the effect of the international agreement on DuPont would be to look at DuPont's stock market valuation. Unfortunately, although in 1986 DuPont produced CFCs for about half of the US market, it represented only 2.2% of DuPont revenues (1.8% in 1984 and 1.7% in 1985), 2% of corporate assets and 0.9% of DuPont's employees (Reinhardt et al. 1989a). It is therefore unlikely that financial markets would have captured much impact. Additionaly, it would difficult to attribute any movement to the regulation of CFCs only and not to other parts of DuPont's business (especially since DuPont was facing other public relations issues related to medical implants of which it supplied the raw material).

6.4 Annex A and Annex B Compounds

In this section, I investigate the effect of Montreal on patents and articles mentioning CFCs, that is the molecules which were being phased out of industrial activities. These molecules are referred to as Annex A compounds because they are listed in the Annex A of the legal text. These molecules include five chlorofluorocarbons and three halons. For chlorofluorocarbons, the agreement imposed a freeze by 1989 and a 50% decrease by 1998 relative to 1986; for halons, only a freeze by 1992 was decided. In 1990, during the London revisions, twelve additional compounds became regulated. They are listed in the Annex B of the agreement and consist of 10 other CFCs plus carbon tetrachloride and methyl chloroform. The negotiated reduction targets for each compound is shown in the Appendix. In what follows, I refer to these two groups of molecules as Annex A and Annex B.

It is difficult to make strong hypotheses about the effect of Montreal for Annex A and Annex B compounds. On one hand, Montreal can be thought of an incentive to no longer pursue any research or innovation that would make use of these molecules in new industrial contexts. But the agreement might also have spurred research efforts to help reduce the ongoing effect of such molecules on the environment as well as innovations to help recycle such components or use them more efficiently. This second effect is particularly likely as the phase-out of such molecules was scheduled to be progressive. As a result, firms were given some time to adapt and could continue using CFCs in their production.

The graphs in Figure B7 plot the yearly number of articles or patents mentioning the names of given molecules included in Annex A and B. We note that most trends are flat, except maybe for Annex A in articles which seem to increase and then decrease. Table B11 presents results from first differences specifications. Results indicate statistically significant mean shifts between before and after 1987, except for Annex B in patents; however these are small in magnitude. In figure B12, the DiD specifications indicate that a positive and statistically significant treatment effect for Annex A in patents and a negative one for Annex B in articles. The magnitudes however are small. For Annex A in patents, the coefficient corresponds to a 18% increase in the number of patents mentioning Annex A compounds. For Annex B in articles, the estimate corresponds to a

28% decrease in the number of articles mentioning Annex B compounds.

Table B13 displays the summary performance of the SCM implementations for Annex A and B in patents and articles. Almost none of the implementations find a significant treatment effect, except for Annex B in articles where a negative treatment effect with 10% significance is found when the whole sample of HAPs is used. These results indicate that Montreal did not trigger a large decrease nor a large increase in the number of patents and articles mentioning Annex A and B compounds. Figure B8 and B9 illustrate these results by displaying the graphs generated by the SCM procedures using the small pool, and unweighted or weighted topic proportions (whichever gave lowest pre-RMPSE). The graphs show that indeed the estimated treatment effect falls well within the distribution of placebo effects, at least for Annex A compounds. For Annex B compounds, the treatment effect tends to be as one of the lowest curves among all the placebos. We note that, in the case of articles, the synthetic control fails to provide a good fit, and so results cannot be trusted.

7 Discussion

The signature of Montreal triggered a series of mechanisms that provided firms and researchers clear incentives to orient their R&D effort towards CFC substitutes. In particular, it had the immediate effect of modifying expectations about future prices and created a worldwide demand for substitutes. Hence, it incentivized profit-seeking firms to bring CFC substitutes to market. Researchers publishing in peer-reviewed journals also redirected their work towards CFC substitutes either incentivized by grants focusing on ozone depletion or due to a shift in their personal research priorities.

I shall recognize here that the treatment is broadly defined as the "ozone regime". As such, it includes the initial diplomatic agreement in Montreal in September 1987, the following country-by-country ratifications[66], and the following amendments (the London, Copehagen and Vienna revisions in 1990, 1992 and 1995 respectively). But it also includes the domestic regulations that were implemented to translate the international agreements into national rules. In the US for

[66]The USA ratified in April 1988; European countries in December 1988.

example, this is done through the Clean Air Act amendment of 1990. The counterfactual therefore represents a world without any of those interventions, and in particular a world without, or with limited, unilateral actions. The inherent challenges of any global public good problem provide a good case for arguying that a world with limited unilateral actions is an apropriate counterfactual. The absence of costly unilateral actions before 1987 is further testimony to those challenges[67].

I shall recognize as well that the counterfactual used here is a world without the "discovery" of the so-called ozone "hole". Indeed, the science of ozone made much progress during the 1980s: in 1985, scientists detected a large depletion of ozone over Antarctica (the "hole") and, importantly, they were able to causally attributed it to CFCs in March 1988 (the "discovery"). The perceived benefits of phasing-out CFCs certainly increased, and it likely contributed to the deepening and widening of the ozone regime in London in 1990 and Copenhagen in 1992: more ambitious reduction targets were then agreed (deepening), while other molecules were added to the list of regulated compounds (widening). What would have technological change looked like in a world without an international agreement but with the discovery of the ozone hole?

The image of the Earth seen from space with a massive hole (artifically colored in blue for the occasion) indeed became world-famous and moved public opinion. But would it have trickled down in terms of individual purchasing decision? Can consumer pressure be a strong enough incentive for firms to transition? There exists a few empirical analysis suggesting this actually happens (Lyon et al. 1999; Popp et al. 2011), but they all deals with local pollutants such as toxis chemical emissions; we could reasonably expect consumer pressure to be less effective for global air pollutant such as ozone. As profit-maximizing entities, firms would have few incentives to incur R&D costs without the guarantee of a large market and without the guarantee that their foreign and domestic competitors do the same. Hence I expect the ozone role to have played any particular role in the dynamics of innovation. I, nonetheless, intend to further this analysis by leveraging variation in the timing of regulation of different molecules. Specifically, two molecules (carbon tetrachloride and methyl chloroform) were recognized as strong ozone-depleting substances but were not included in the 1987 agreement. Instead, they became regulated in 1990. I intend to analyze trends in science

[67]The aerosol bans in the 70s and 80s were not costly because physical and chemical substitutes existed; for example, roll-on deodoratn instead of spray deodorants.

and innovation for substitutes to carbon tetrachloride and methyl chloroform to infer whether the ozone hole "discovery" might have directly affected firms' behavior.

8 Conclusion

Tackling environmental problems often relies on developing and diffusing new technologies. It is, therefore, important to better understand the drivers of technological change. In this paper, I document that the Montreal Protocol, and its following amendments, led to the development of CFCs substitutes. This empirical evidence goes against the often-heard narrative that alternatives technologies were readily available before the treaty. Instead, the treatment effect that I estimate in this paper tells a story where almost all of the science and innovation on CFC substitutes was triggered by the post-Montreal regime. The magnitude of the effect is even consistent with what has been described as a "burst of industrial creativity" (Meadows et al. 1992).

For sociologist Reiner Grundmann , the idea that CFCs substitutes were already available is "the most pervasive and most widespread myth surrounding the Montreal Protocol" (Grundmann 1998). He traced its origin to the fierce opposition between Americans and Europeans during the Concorde controversy when the French and the British hoped to conquer the world with supersonic jets. In 1974, Americans denied the authorization to land the aircraft on U.S. soil, protesting that the pollution emitted by the plane's engine was a serious threat to the ozone layer. The claims were eventually dismissed because it was shown that the atmospheric reactions did not occur at the Concorde flying altitudes (Benedick 2009, p. 33). But the incident certainly left an aftertaste of distrust for Europeans, and maybe primed them to assume that, when dealing with Americans, environmental issues are but a disguise for commercial interests. When CFC regulations came onto the international agenda, it probably was only a small step for European manufacturers to assume that the pro-regulatory stance displayed by Americans was motivated by the existence of secret substitute which, with an international ban on CFCs, would have allowed US manufacturers to gain market shares.

9 References

Abadie, Alberto and Javier Gardeazabal (2003). "The Economic Costs of Conflict: A Case Study of the Basque Country." In: *The American Economic Review* 93.1, pp. 113–132.

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller (2010). "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program." In: *Journal of the American Statistical Association* 105.490, pp. 493–505.

— (2015). "Comparative Politics and the Synthetic Control Method." In: *American Journal of Political Science* 59.2, pp. 495–510.

Acemoglu, Daron (1998). "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality." In: *The Quarterly Journal of Economics* 113.4, pp. 1055–1089.

— (2002). "Directed Technical Change." English. In: *The Review of Economic Studies* 69.4, pp. 781–809. ISSN: 00346527.

Aghion, Philippe et al. (2016). "Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry." In: *Journal of Political Economy* 124.1, pp. 1–51.

Aichele, Rahel and Gabriel Felbermayr (2011). "Kyoto and the Carbon Footprint of Nations." In: *Journal of Environmental Economics and Management* 63.3, pp. 336–354.

Ambec, Stefan et al. (2013). "The Porter Hypothesis at 20: Can Environmental Regulation Enhance Innovation and Competitiveness?" In: *Review of Environmental Economics and Policy* 7.1, p. 2.

Athey, Susan and Guido Imbens (2016). "The State of Applied Econometrics-causality and Policy Evaluation." In: *Arxiv Preprint Arxiv:1607.00699*.

Barrett, Scott (1994). "Self-enforcing International Environmental Agreements." In: *Oxford Economic Papers*, pp. 878–894.

— (1999). "Montreal Versus Kyoto: International Cooperation and the Global Environment." In: *Global public goods: international cooperation in the 21st century*. Oxford University Press.

— (2003). *Environment and Statecraft: The Strategy of Environmental Treaty-making: The Strategy of Environmental Treaty-making*. Oxford University Press.

Benedick, Richard Elliot (2009). *Ozone Diplomacy: New Directions in Safeguarding the Planet*. Harvard University Press.

Beron, Kurt J, James C Murdoch, and Wim PM Vijverberg (2003). "Why Cooperate? Public Goods, Economic Power, and the Montreal Protocol." In: *The Review of Economics and Statistics* 85.2, pp. 286–297.

Blei, David M (2012). "Probabilistic Topic Models." In: *Communications of the Acm* 55.4, pp. 77–84.

Blei, David M and John D Lafferty (2006). "Dynamic Topic Models." In: pp. 113–120.

— (2009). "Topic Models." In: *Text Mining: Classification, Clustering, and Applications* 10.71, p. 34.

Brunel, Claire (2015). "Green Innovation and Green Manufacturing: Links between Environmental Policies, Innovation and Production." Working paper, 2015.

Brunnermeier, Smita B and Mark A Cohen (2003). "Determinants of Environmental Innovation in U.S. Manufacturing Industries." In: *Journal of Environmental Economics and Management* 45.2, pp. 278–293.

Calel, Raphael and Antoine Dechezleprêtre (2016). "Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market." In: *Review of Economics and Statistics* 98.1, pp. 173–191.

Dasgupta, Partha and Paul A. David (1994). "Toward a New Economics of Science." In: *Research Policy* 23.5. Special Issue in Honor of Nathan Rosenberg, pp. 487 –521. ISSN: 0048-7333.

Dekker, Thijs et al. (2012). "Inciting Protocols." In: *Journal of Environmental Economics and Management* 64.1, pp. 45–67.

Falkner, Robert (2005). "The Business of Ozone Layer Protection: Corporate Power in Regime Evolution." In: *The Business of Global Environmental Governance*. MIT Press.

Finus, Michael and Sigve Tjøtta (2003). "The Oslo Protocol on Sulfur Reduction: The Great Leap Forward?" In: *Journal of Public Economics* 87.9, pp. 2031–2048.

Glynn, Steven (2002). "Constructing a Selection Environment: Competing Expectations for CFC Alternatives." In: *Research Policy* 31.6, pp. 935–946.

Gonzalez, Marco, Kristen N Taddonio, and Nancy J Sherman (2015). "The Montreal Protocol: How Today's Successes Offer a Pathway to the Future." In: *Journal of Environmental Studies and Sciences* 5.2, pp. 122–129.

Greenhalgh, Christine and Mark Rogers (2010). *Innovation, Intellectual Property, and Economic Growth*. Princeton University Press.

Grundmann, Reiner (1998). "The Strange Success of the Montreal Protocol-why Reductionist Accounts Fail." In: *International Environmental Affairs* 10.3, pp. 197–220.

Hall, Bronwyn H and Adam B Jaffe (2012). "Measuring Science, Technology, and Innovation: A Review." In: *Report Prepared for the Panel on Developing Science, Technology, and Innovation Indicators for the Future*.

Heal, Geoffrey (2016). *Endangered Economies: How the Neglect of Nature Threatens Our Prosperity*. Ed. by Columbia University Press. Columbia University Press.

Hegglin, Michaela I., Stephen A. Montzka David. W. Fahey Mack McFarland, and Eric R. Nash (2015). *Twenty Questions and Answers about the Ozone Layer: 2014 Update. Scientific Assessment of Ozone Depletion: 2014, 84 Pp.,* tech. rep. World Meteorological Organization, Geneva, Switzerland,

Henderson, Rebecca, Adam B Jaffe, and Manuel Trajtenberg (1998). "Universities As a Source of Commercial Technology: A Detailed Analysis of University Patenting, 1965-1988." In: *Rev Econ Stat* 80.1, pp. 119–127.

Hicks, John R (1932). "The Theory of Wages." In: *Et Seq*, p. 54.

Iaria, Alessandro and Fabian Waldinger (2015). "International Knowledge Flows: Evidence from the Collapse of International Science in the Wake of WWI."

Jaffe, Adam B and Karen Palmer (1997). "Environmental Regulation and Innovation: A Panel Data Study." In: *Review of Economics and Statistics* 79.4, pp. 610–619.

Jaffe, Adam B, Richard G Newell, and Robert N Stavins (2002). "Environmental Policy and Technological Change." In: *Environmental and Resource Economics* 22.1, pp. 41–70.

Johnstone, Nick, Ivan Haščič, and David Popp (2010). "Renewable Energy Policies and Technological Innovation: Evidence Based on Patent Counts." In: *Environmental and Resource Economics* 45.1, pp. 133–155.

Johnstone, Nick et al. (2012). "Environmental Policy Stringency and Technological Innovation: Evidence from Survey Data and Patent Counts." In: *Applied Economics* 44.17, pp. 2157–2170.

Kay, Luciano et al. (2014). "Patent Overlay Mapping: Visualizing Technological Distance." In: *Journal of the Association for Information Science and Technology* 65.12, pp. 2432–2443.

Kellenberg, Derek and Arik Levinson (2014). "Waste of Effort? International Environmental Agreements." In: *Journal of the Association of Environmental and Resource Economists* 1.1/2, pp. 135–169.

Lanjouw, Jean Olson and Ashoka Mody (1996). "Innovation and the International Diffusion of Environmentally Responsive Technology." In: *Research Policy* 25.4, pp. 549–571.

Le Prestre, Philippe G, John D Reid, and E Thomas Morehouse Jr (1998). *Protecting the Ozone Layer: Lessons, Models, and Prospects*. Springer Science & Business Media.

Lyon, Thomas P. and John W. Maxwell (1999). "'Voluntary' Approaches to Environmental Regulation: A Survey." In: *Ssrn Electronic Journal*. ISSN: 1556-5068. DOI: 10.2139/ssrn.147888. URL: http://www.ssrn.com/abstract=147888.

Meadows, Donella H, Dennis L Meadows, Jørgen Randers, et al. (1992). *Beyond the Limits: Global Collapse or a Sustainable Future.* Earthscan Publications Ltd.

Miller, Alan S and Irving M Mintzer (1986). "Sky Is the Limit; Strategies for Protecting the Ozone Layer." World Resources Institute, Research Report, 1986.

Molina, Mario J and F Sherwood Rowland (1974). "Stratospheric Sink for Chlorofluoromethanes: Chlorine Atom-catalysed Destruction of Ozone." In: *Nature* 249.28, pp. 810–812.

Moretti, Edward C (2002). "Reduce VOC and HAP emissions." In: *Chem. Eng. Prog.* 98.6, pp. 30–40.

Mulder, Karel F (2005). "Innovation by Disaster: The Ozone Catastrophe As Experiment of Forced Innovation." In: *International Journal of Environment and Sustainable Development* 4.1, pp. 88–103.

Murdoch, James C and Todd Sandler (1997). "The Voluntary Provision of a Pure Public Good: The Case of Reduced CFC Emissions and the Montreal Protocol." In: *Journal of Public Economics* 63.3, pp. 331–349.

— (2009). "The Voluntary Provision of a Pure Public Good and the Montreal Protocol: Behavioral and Data Concerns." In: *Oxford Economic Papers* 61.1, pp. 197–200.

Nesta, Lionel, Francesco Vona, and Francesco Nicolli (2014). "Environmental Policies, Competition and Innovation in Renewable Energy." In: *Journal of Environmental Economics and Management* 67.3, pp. 396–411.

Newell, Richard G., Adam B. Jaffe, and Robert N. Stavins (1999). "The Induced Innovation Hypothesis and Energy-saving Technological Change." In: *Quarterly Journal of Economics* 114.458, pp. 907–940.

Parson, Edward A (2003). *Protecting the Ozone Layer: Science and Strategy.* Oxford University Press.

Popp, David (2002). "Induced Innovation and Energy Prices." In: *American Economic Review*, pp. 160–180.

— (2005). "Lessons from Patents: Using Patents to Measure Technological Change in Environmental Models." In: *Ecological Economics* 54.2, pp. 209–226.

— (2006). "They Don't Invent Them like They Used to: An Examination of Energy Patent Citations Over Time." In: *Economics of Innovation and New Technology* 15.8, pp. 753–776.

— (2010). "Exploring Links between Innovation and Diffusion: Adoption of NOX Control Technologies at U.S. Coal-fired Power Plants." In: *Environmental and Resource Economics* 45.3, pp. 319–352.

Popp, David, Richard G Newell, and Adam B Jaffe (2010). "Chapter 21: Energy, the Environment, and Technological Change." In: *Handbook of the Economics of Innovation* 2, pp. 873–937.

Popp, David, Tamara Hafner, and Nick Johnstone (2011). "Environmental policy vs. public pressure: Innovation and diffusion of alternative bleaching technologies in the pulp industry." In: *Research Policy* 40.9, pp. 1253–1268. ISSN: 0048-7333. DOI: 10.1016/J.RESPOL.2011.05.018. URL: https://www.sciencedirect.com/science/article/pii/S0048733311000990.

Porter, Michael E. (1991). "America's Green Strategy." In: *Scientific American* 264.4, p. 168.

Porter, Michael E and Claas Van Der Linde (1995a). "Green and Competitive: Ending the Stalemate." In: *Harvard Business Review* 73.5, pp. 120–134.

— (1995b). "Toward a New Conception of the Environment-competitiveness Relationship." In: *The Journal of Economic Perspectives*, pp. 97–118.

Price, Derek J. De Solla (1965). "Networks of Scientific Papers." In: *Science* 149.3683, pp. 510–515.

— (1986). *Little Science, Big Science... and Beyond*. Columbia University Press New York.

Puller, Steven L (2006). "The Strategic Use of Innovation to Influence Regulatory Standards." In: *Journal of Environmental Economics and Management* 52.3, pp. 690–706.

Redner, Sidney (2004). "Citation Statistics from More Than a Century of Physical Review." In: *Arxiv Preprint Physics/0407137*.

Reinhardt, Forest and Richard HK Vietor (1989a). "Du Pont Freon Products Division (a)." In: *Harvard Business School Case 389-111*. Revised March 1995, pp. 261–286.

— (1989b). "Du Pont Freon Products Division (b)." In: *Managing Environmental Issues: A Casebook*, pp. 261–286.

Roberts, Margaret E, Brandon M Stewart, and Dustin Tingley (2014). "Stm: R Package for Structural Topic Models." In: *R Package* 1, p. 12.

Roberts, Margaret, Brandon Stewart, and Dustin Tingley (2016). "Navigating the Local Modes of Big Data: The Case of Topic Models." In: *Data Analytics in Social Science, Government, and Industry*. Cambridge University Press, New York.

Smith, Brigitte (1998). "Ethics of Du Pontś CFC Strategy 1975-1995." In: *Journal of Business Ethics* 17.5, pp. 557–568.

Sunstein, Cass R (2007). "Of Montreal and Kyoto: A Tale of Two Protocols." In: *Harv. Envtl. L. Rev.* 31, p. 1.

Taddonio, Kristen, K Madhava Sarma, and Stephen O Andersen (2012). *Technology Transfer for the Ozone Layer: Lessons for Climate Change*. Routledge.

Thompson, Peter and Melanie Fox-Kean (2005). "Patent Citations and the Geography of Knowledge Spillovers: A Reassessment." In: *Am Econ Rev*, pp. 450–460.

Trajtenberg, Manuel, Rebecca Henderson, and Adam Jaffe (1997). "University Versus Corporate Patents: A Window on the Basicness of Invention." In: *Economics of Innovation and New Technology* 5.1, pp. 19–50.

Vollebergh, Herman RJ (2007). "Differential Impact of Environmental Policy Instruments on Technological Change: A Review of the Empirical Literature."

Wagner, Ulrich J (2009). "The Voluntary Provision of a Pure Public Good? Another Look at Cfc Emissions and the Montreal Protocol." In: *Oxford Economic Papers* 61.1, pp. 183–196.

— (2016). "Estimating Strategic Models of International Treaty Formation." In: *The Review of Economic Studies* 83.4, pp. 1741–1778.

Wang, Dashun, Chaoming Song, and Albert-László Barabási (2013). "Quantifying Long-term Scientific Impact." In: *Science* 342.6154, pp. 127–132.

Williams, Heidi L (2013). "Intellectual Property Rights and Innovation: Evidence from the Human Genome." In: *Journal of Political Economy* 121.1, pp. 1–27.

Williams, Heidi (2017). "Patents and Research Investments: What Inventions Are We Missing?" In: *Annual Review of Economics* 9.1.

Figures and Tables

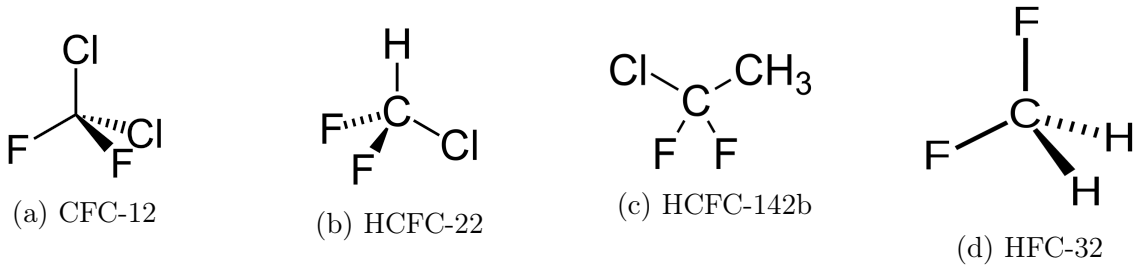

(a) CFC-12 (b) HCFC-22 (c) HCFC-142b (d) HFC-32

Figure 1: Typical Molecular Structure of CFCs, HCFCs and HFCs

Note: CFC stands for chlorofluorocarbon, i.e. a molecule entirely made of carbon, chlorine and florine atoms. When one chlorine atom in CFC-12 is substituted with a hydrogen, it becomes HCFC-22. If substituted with a methyl group, we obtain HCFC-142b. Here "HCFC" stands for hydro-chlorofluorocarbons. When all chlorine atoms are substituted with hydrogens, the compounds are then known as HFC, or hydro-fluorocarbons. For example, when the two chlorine atoms in CFC-12 are replaced by hydrogens, we get HFC-32.

Table 1: Ten most common patent codes for patents mentioning CFC substitutes

ICL	Count	Description
C07C	501	Acyclic or carbocyclic compounds
C08G	351	Compounds of unknown constitution
C08J	269	General processes of compounding
C09K	183	Materials for applications not otherwise provided for
A61K	156	Preparations for medical, dental, or toilet purposes
C10M	106	Lubricating compositions
F25B	90	Refrigeration machines, plants, or systems; heat pump systems
C08F	67	Macromolecular compounds obtained by reactions only involving carbon-to-carbon unsaturated bonds
C11D	64	Detergent compositions
C07D	63	Heterocyclic compounds

Note: The table displays the most frequent codes associated to patents mentioning CFC substitutes throughout the period 1976 to 2000. Most codes belong to the C class ("Chemistry, Metallurgy"). The subclasses "C07" and "C08" refer in particular to the preparation (e.g., purification, separation or stabilisation) of organic compounds, and as such they are associated to any patent related to compounds containing carbon and halogens with or without hydrogen (e.g., C07C 19/00: Acyclic saturated compounds containing halogen atoms). Subclass 'C08G' is used to classify any preparation that uses fermentation or enzyme-using processes. Only patents with at least 3 molecule occurrences are kept in the sample.

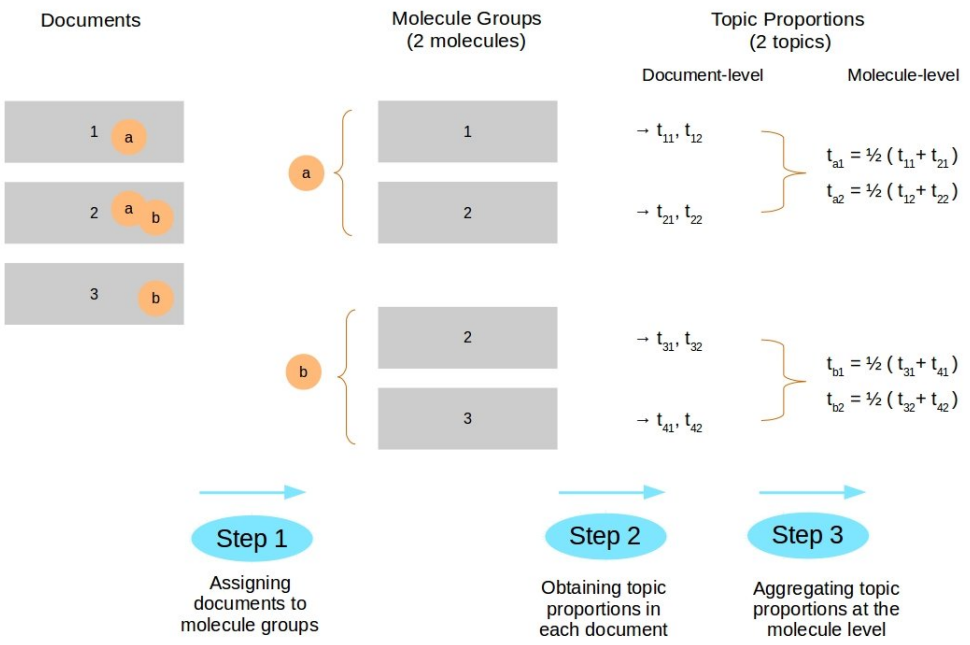

Figure 2: Schematic explanation of the methodology

Note: Suppose there are three documents: document 1 and 2 mention molecule 'a' while document 2 and 3 mention molecule 'b'. In step 1, I aggregate documents according to their molecule group. I follow a basic rule that assign any document with at least one mention of a molecule to that molecule's group. In step 2, I use topic modeling to obtain the proportions of topics in each document. t_i, j stands for the proportion of topic j in document i. Finally, in step 3, I create a topic proportion at the molecule level by averaging over all the documents that mention the molecule of interest.

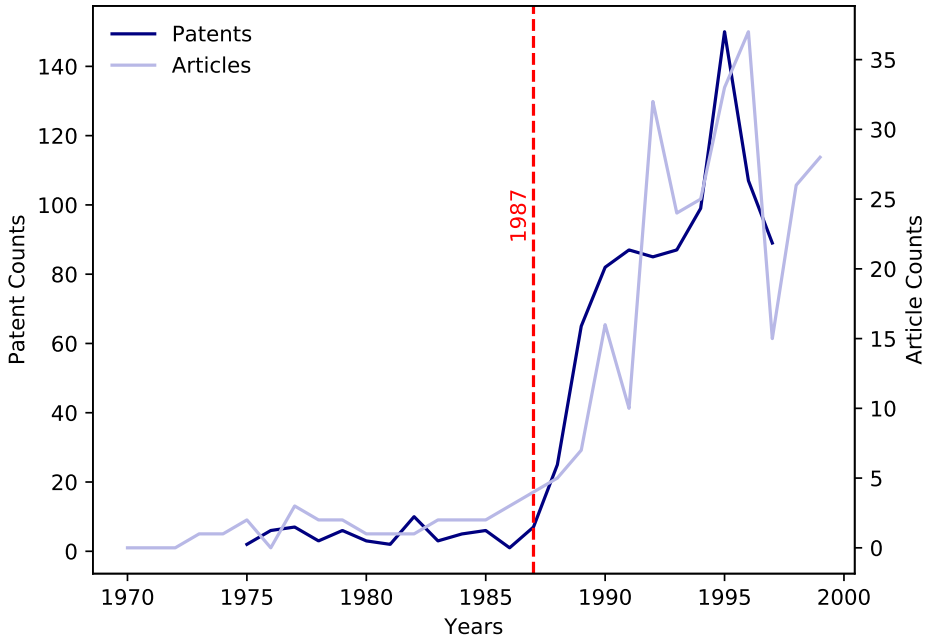

Figure 3: Counts of patents and articles on CFC substitutes

Note: The graph plot the yearly number of articles or patents mentioning the names of any of the 14 CFC substitutes. This is not the average count of the 14 different CFC substitutes but the total count of documents mentioning any one of the 14 CFC substitutes. We note a clear increase for both patents and articles after 1987, the year Montreal was signed. For patents, the graph shows any patent *granted* (as opposed to patent applications) between 1976 and 1999. The year on the x-axis, however, corresponds to the application date. There is on average a two-year delay between patent application and grant. For articles, the year on the x-axis corresponds to the year the article was published in the academic journal.

Table 2: First differences for CFC substitutes

(a) Patents

	(1)	(2)
Post 1987	1.840***	1.268***
	(0.077)	(0.165)
Post 1987 x Years		0.095***
		(0.026)
Years		0.000
		(0.012)
Molecule FE	Yes	Yes
R-squared	0.789	0.809
Observations	322	322

Standard errors in parentheses
Dependent variable: Log count of patents
Years are relative to 1987.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

(b) Articles

	(1)	(2)
Post 1987	1.053***	0.344***
	(0.067)	(0.119)
Post 1987 x Years		0.083***
		(0.016)
Years		0.011*
		(0.006)
Molecule FE	Yes	Yes
R-squared	0.649	0.693
Observations	420	420

Standard errors in parentheses
Dependent variable: Log count of articles
Years are relative to 1987.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The tables present regression results for first difference specifications. Model 1 confirms that there is a significant and positive mean shift after 1987 in the number of patents and articles mentioning CFC substitutes. Model 2 indicates that the change can also be modeled as a trend break. The coefficient for 'Years' indicates that there is a small but statistically significant positive underlying trend for both articles and patents.

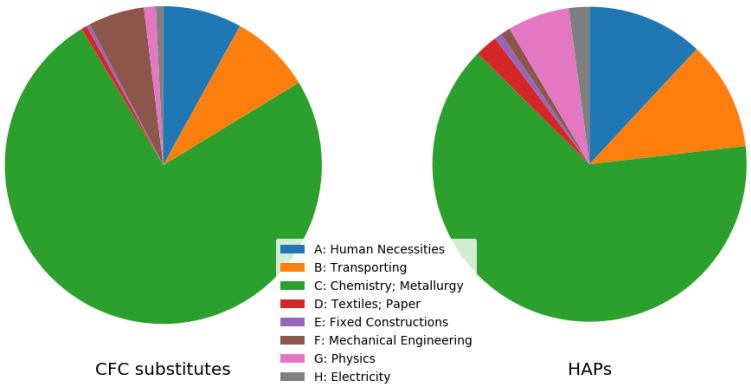

Figure 4: Top level patent codes for CFC substitutes and HAPs

Note: The figure shows that, overall, patents mentioning CFC substitutes and HAPs fall into similar top-level codes. HAPs are a group of 171 molecules that have no relationship to ozone and that are used for diverse industrial applications. The figure indicates the two groups of molecules present important similarities which motivates the use of HAPs as control molecules to estimate the causal effect of the post-Montreal regime. The patent codes are from the international patent classification.

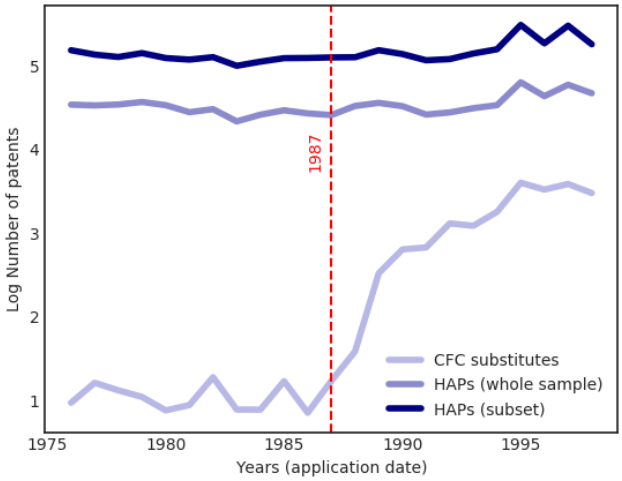

(a) Patents

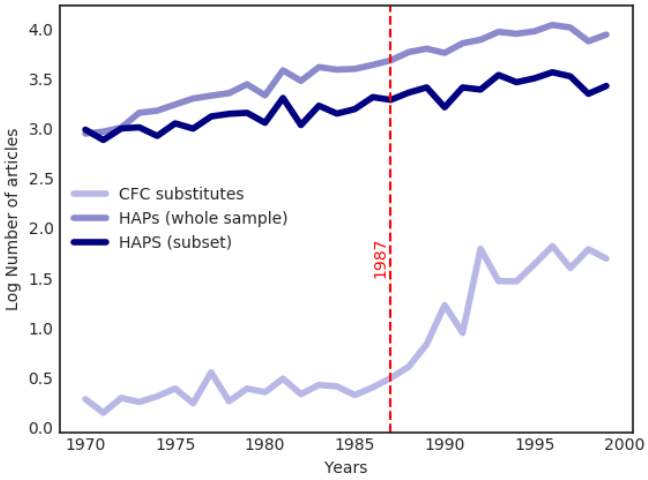

(b) Articles

Figure 5: Pre-trends in log counts of documents mentioning CFC substitutes and HAPs

Note: The graphs display the pre-trends for the treated group (CFC substitutes) and two possible control groups. The first uses the whole sample of HAPs (that is 171 molecules). In this case, pre-trends look somewhat similar. Pre-trends are, however, closer when using a smaller subset. Specifically, the second control group, shown as "subset" on the graphs, uses a smaller number of HAPs: the 40 HAPs with pre-trends closest to the average CFC substitutes.

Table 3: Difference-in-differences for CFC substitutes

(a) Patents

	(1)	(2)
Post 1987 x Substitutes	1.858***	1.078***
	(0.072)	(0.169)
Post 1987 x Substitutes x Years		0.095***
		(0.026)
Substitutes x Years		0.018
		(0.012)
Years		-0.018***
		(0.004)
Post 1987		0.190***
		(0.036)
Year FE	Yes	No
Molecule FE	Yes	Yes
R-squared	0.974	0.967
Observations	1288	1288

Standard errors in parentheses
Dependent variable: Log count of patents
Years are relative to 1987.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

(b) Articles

	(1)	(2)
Post 1987 x Substitutes	0.727***	0.291**
	(0.068)	(0.126)
Post 1987 x Substitutes x Years		0.083***
		(0.016)
Post 1987		0.052
		(0.041)
Substitutes x Years		-0.007
		(0.006)
Years		0.018***
		(0.002)
Year FE	Yes	No
Molecule FE	Yes	Yes
R-squared	0.947	0.947
Observations	1680	1680

Standard errors in parentheses
Dependent variable: Log count of articles
Years are relative to 1987.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note:

The tables present regression results for the difference-in-differences specifications. Model 1 corresponds to the main DiD specification. It includes year and molecule fixed effects. The binary variable 'Post 1987 x Substitutes' equals 1 for observations belonging to the group CFC substitutes and after 1987. For patents, the coefficient is smaller than the coefficient in the simple difference, but remains significant and large, corresponding to close to a 550% increase. For articles, the coefficient is smaller than the coefficient in the simple difference, but remains significant and large, corresponding to more than a 95% increase. Model 2 presents a trend-break specification. It shows that the log number of patents mentioning CFC substitutes increases with the years after 1987 by 0.22 more than the control group. Similarly, the log number of articles mentioning CFC substitutes increases with the years after 1987 by 0.10 more than the control group.

(a) Patents

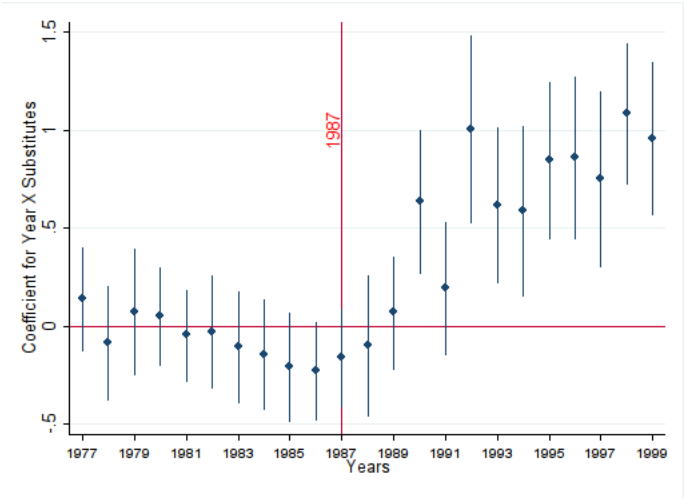

(b) Articles

Figure 6: Difference-in-differences treatment effects by year

Note: The graphs display the DiD coefficients for each year. We note that, in patents, the treatment effect is statistically significant, yet small, as early as 1988. For articles, the treatment effect is first statistically significant in 1990.

Table 4: Balance table between CFC substitutes and HAPs

(a) Patents

	HAPs	CFC substitutes	Difference	T-stat
Number of patents	465.70	16.37	449.34***	(13.12)
Log number of patents	4.91	1.93	2.98***	(24.19)
Weighted mean proportion of topic 1	0.19	0.16	0.04***	(9.24)
Weighted mean proportion of topic 2	0.10	0.09	0.02***	(7.14)
Weighted mean proportion of topic 3	0.41	0.33	0.08***	(10.22)
Weighted mean proportion of topic 4	0.18	0.13	0.04***	(11.59)
Weighted mean proportion of topic 5	0.07	0.07	0.00	(1.67)

(b) Articles

	HAPs	CFC substitutes	Difference	T-stat
Number of articles	159.92	3.28	156.64***	(6.46)
Log number of articles	3.23	0.78	2.46***	(26.21)
Weighted mean proportion of topic 1	0.19	0.11	0.09***	(21.10)
Weighted mean proportion of topic 2	0.30	0.15	0.15***	(26.70)
Weighted mean proportion of topic 3	0.12	0.05	0.06***	(24.54)
Weighted mean proportion of topic 4	0.19	0.09	0.09***	(24.76)
Weighted mean proportion of topic 5	0.16	0.07	0.10***	(28.34)

Note: The table shows the mean of the outcome variable (in log and in level) and the topic proportions for patents and articles on CFC substitutes and HAPs. We see that the two groups have very different average counts. Mean topic proportions are also statistically different across the two groups.

Table 5: Robustness checks difference-in-differences

(a) Patents

	(1)	(2)	(3)	(4) Unweighted	(5) Weighted
Post 1987 x Substitutes	1.865***	0.904***	0.845***	1.583***	1.210***
	(0.072)	(0.092)	(0.092)	(0.076)	(0.096)
Log Count (lag 1)		0.559***	0.371***		
		(0.039)	(0.049)		
Log Count (lag 2)			0.270***		
			(0.047)		
Mean proportion of topic 1				2.088***	1.128*
				(0.779)	(0.619)
Mean proportion of topic 2				-0.199	0.124
				(0.584)	(0.554)
Mean proportion of topic 3				1.194***	1.212***
				(0.391)	(0.342)
Mean proportion of topic 4				0.321	0.868
				(0.610)	(0.552)
Mean proportion of topic 5				-0.424	0.495
				(1.132)	(0.789)
Year FE	Yes	Yes	Yes	Yes	Yes
Molecule FE	Yes	Yes	Yes	Yes	Yes
R-squared	0.976	0.985	0.986	0.981	0.984
Observations	1288	1232	1176	1285	914

(b) Articles

	(1)	(2)	(3)	(4) Unweighted	(5) Weighted
Post 1987 x Substitutes	0.727***	0.508***	0.421***	0.266***	0.274***
	(0.068)	(0.063)	(0.060)	(0.058)	(0.058)
Log Count (lag 1)		0.348***	0.245***		
		(0.031)	(0.030)		
Log Count (lag 2)			0.297***		
			(0.031)		
Mean proportion of topic 1				1.513***	1.018***
				(0.432)	(0.363)
Mean proportion of topic 2				0.833**	1.096***
				(0.346)	(0.278)
Mean proportion of topic 3				1.108**	1.040**
				(0.481)	(0.409)
Mean proportion of topic 4				1.168***	1.172***
				(0.443)	(0.354)
Mean proportion of topic 5				0.939***	1.094***
				(0.359)	(0.296)
Year FE	Yes	Yes	Yes	Yes	Yes
Molecule FE	Yes	Yes	Yes	Yes	Yes
R-squared	0.947	0.953	0.956	0.961	0.961
Observations	1680	1624	1568	1680	1680

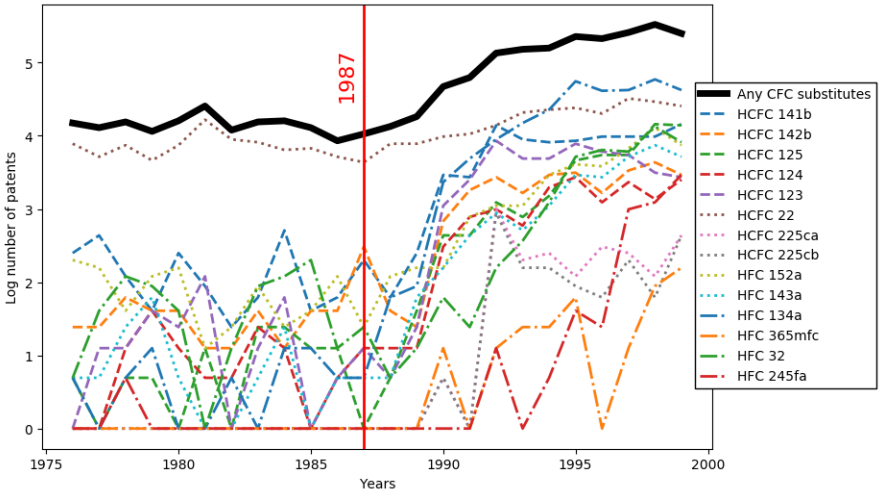

Figure 7: Patent counts in log for CFC substitute, individually and aggregated

Note: The graph illustrates the difference between considering the 14 molecules independently and considering them as one treated molecule. The thick line called "Substitutes (aggregated)" corresponds to the number of patents mentioning any of the 14 CFC substitutes. It is equivalent to considering the 14 compounds as one and only one molecule. I implement the synthetic control method on this "aggregated CFC substitute" because my objective is to estimate the effect of Montreal on research and innovation on any of the CFC subsitutes as opposed to any one in particular. It should be noted here that since the names of different CFC substitutes often appear simultaneously in the same documents, the individual time series of each CFC substitute are not independant from each other.

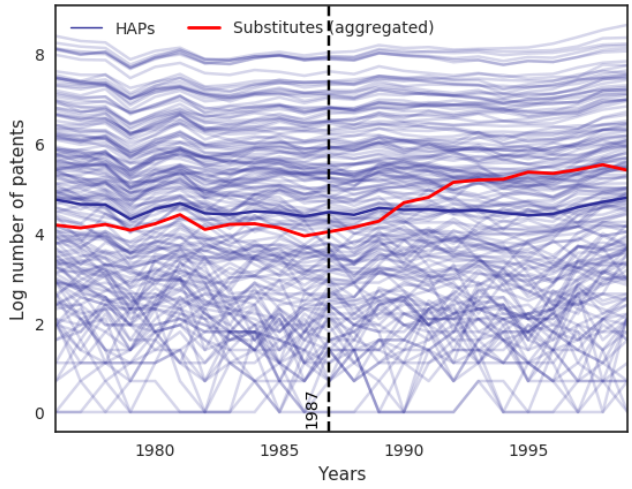

Figure 8: Patent counts in log for each HAP and for the aggregated CFC substitutes

Note: The graph illustrates the heterogeneity of HAP molecules. The thin lines correspond to the trends for each individual HAP while the thick HAP line corresponds to the mean counts for HAPs. We see that HAPs are a diverse group of molecules. In particular, some of them have log counts much higher than the aggregated CFC substitutes. The synthetic control method will allow to construct a better control group by using only the HAPs most similar to CFC substitutes.

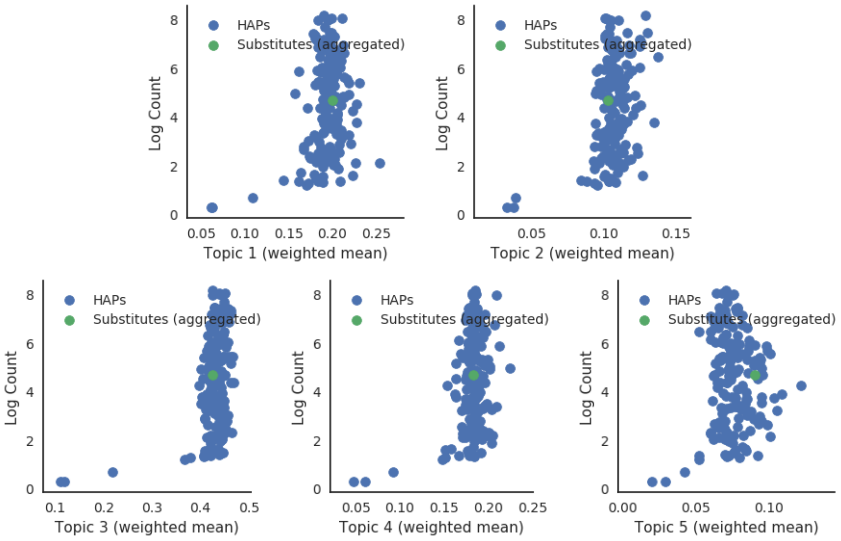

(a) Patents

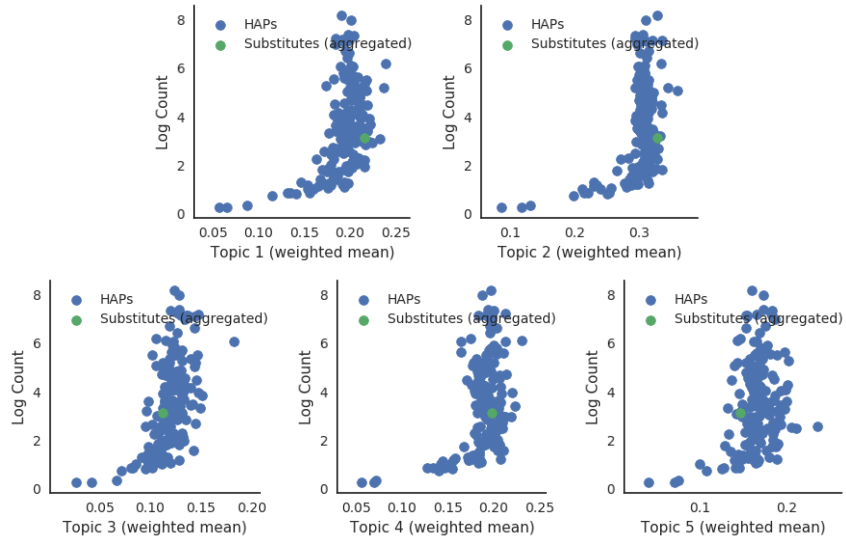

(b) Articles

Figure 9: Scatterplot of topics proportion and log count.

Note: The graphs illustrate the usefulness of topic proportions in the SCM. The scatter plots indicate that there are some clear outlier molecules: molecules with semantic contexts far from CFC substitutes. Implementing the SCM with topic proportions therefore provides a way of avoiding such molecules contribute to constructing a comparison unit. I implement the SCM in two different ways. First, I use the entire sample of HAPs as donor pool (168 units). Second, I create a "small" donor pool containing only the 20 HAPs that are closest to the aggregated CFC substitutes in terms of log counts and topic proportions. Implementing the SCM on a smaller donor pool allows for reducing the risk of overfitting and interpolation bias.

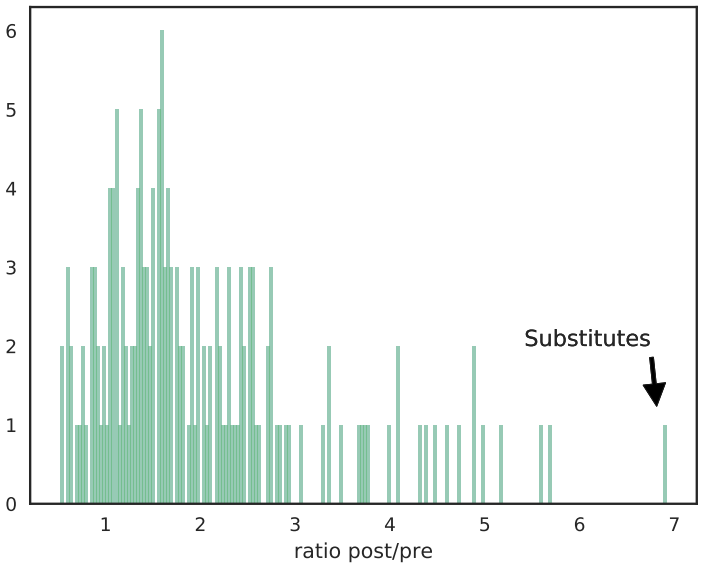

Figure 10: Distribution of post / pre RMSPE ratios for placebos for CFC substitutes

Note: The figure illustrates the ifnerence procedure for the SCM. The graph displays the distribution of post-RMPSE over pre-RMPSE for all placebo units. The figure shows that the ratio for CFC substitutes is clearly greater than all of the 168 other units. Hence the p-value in this case is 1/168.

Table 6: SCM results for CFC substitutes

(a) Patents

Topic Means	Donor Pool	Pre RMSPE	p-value	ATE	Year
unweighted	whole sample	0.12	< 0.01	0.64	1990
weighted	whole sample	0.14	< 0.01	0.83	1990
unweighted	small pool	0.18	< 0.01	0.89	1990
weighted	small pool	0.32	0.11	0.99	1990

(b) Articles

Topic Means	Donor Pool	Pre RMSPE	p-value	ATE	Year
weighted	small pool	0.21	< 0.01	1.06	1990
unweighted	small pool	0.24	< 0.01	1.01	1992
weighted	whole sample	0.34	0.02	1.03	1990
unweighted	whole sample	0.35	< 0.01	1.19	1990

Note: The tables present the results of the main SCM specifications for patents and articles using either weighted or unweighted means of topic proportions and using either the whole sample and small pool of HAPs. The prefered specification to report the average treatment effect (ATE) uses the small pool of HAPs because it minimizes the risk of interpolation biaises and overfitting. It also uses weighted means of topic proportions because it yields a lower pre-RMSPE (pretreatment root mean squared prediction error) indicating that it provides a better counterfactual. For patents, the ATE of the prefered specification is 0.89, that is a 140% increase in patents compared to synthetic control. This corresponds to about 120 patents per year from 1988 to 2000. For articles, the ATE is 1.06, that is a 190% increase in patents compared to synthetic control. This corresponds to about 40 patents per year from 1988 to 2000. "Topic Means" indicates the procedure for aggregating the topic proportions at the molecule level. If "weighted", the calculated proportion of topic j for molecule i is the mean proportion of topic j across all documents mentioning molecule i, weighted by the number of times the molecule appears in the document. "Donor Pool" indicates what sample of HAPs is used in the SCM procedure. For "small pool", the sample of HAPs used corresponds to the twenty HAPs most similar to the treated unit in terms of counts and topic proportions before 1987. "Year" indicates the first year when the treatment effect is significant at the 10% level.

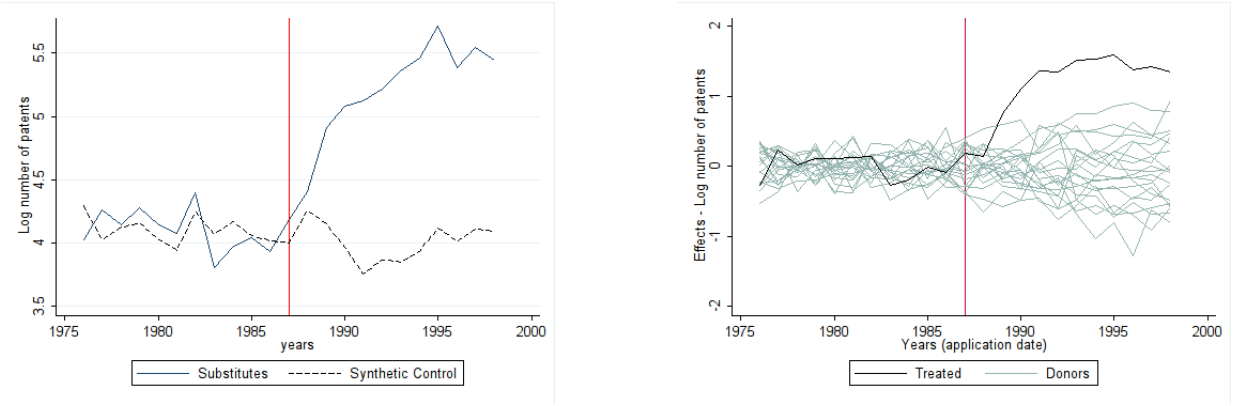

(a) Patents: raw effect (left) and placebo tests (right)

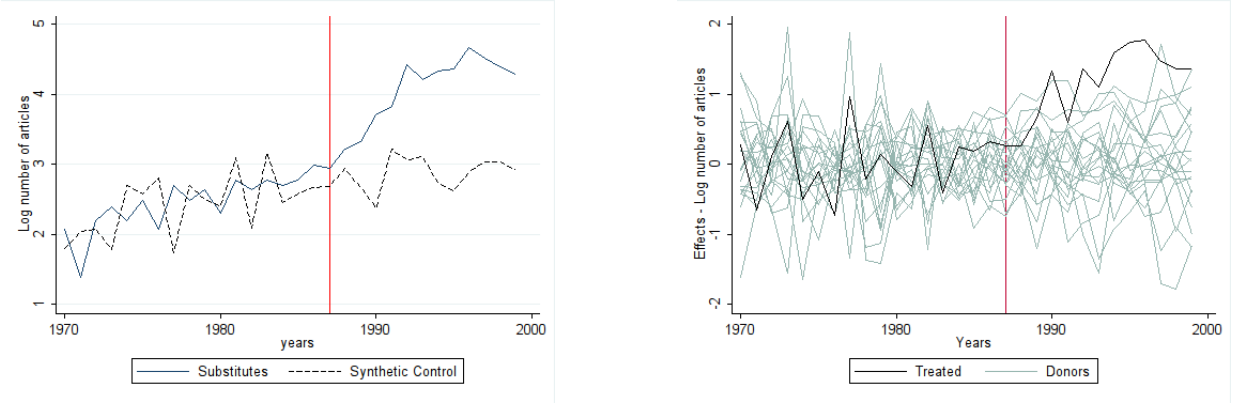

(b) Articles: raw effect (left) and placebo tests (right)

Figure 11: SCM graphs for CFC substitutes

Note: The graphs correspond to SCM implementation for the prefered specification, that is when the SCM is implemented with log counts using the small pool of HAPs and weighted means of topic proportions. The graphs on the left-hand side show the raw effect, that is the observed time series of the treated group along with the time series of the constructed control. On the right-hand sides are shown the placebo tests, the non-parametric tests to evaluate the significance of the results; black lines show the relative effect on the treated group relative to the control group, while each gray line is a placebo test performed on an unit drawn from the donor pool. The effect on CFC substitutes appears large and significant for both patents and articles. Placebo tests confirm that the effect is significant starting a few years after 1987; indeed the black line rises above most other lines as from 1990. This might correspond to a natural lag time between the redirection of research activities towards CFC substitutes and the publication or patenting of such work.

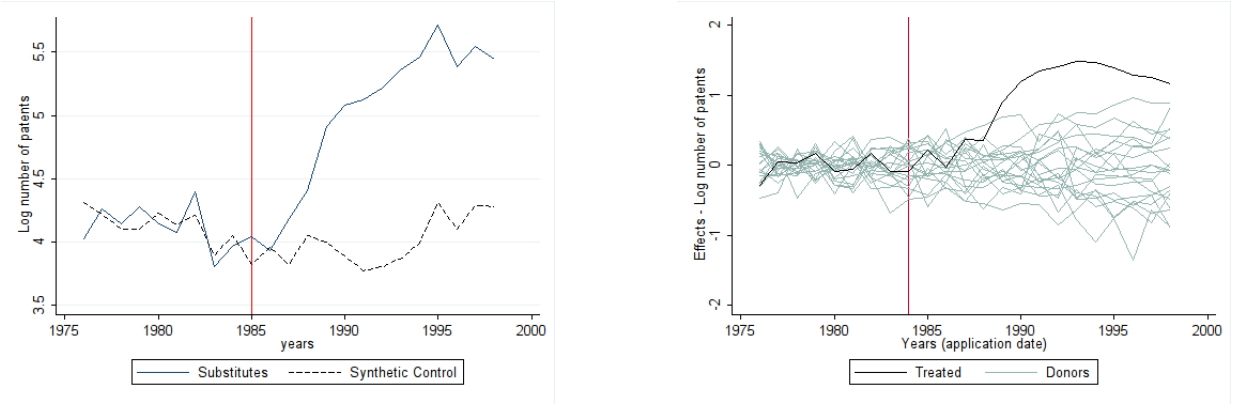

(a) Patents: raw effect (left) and placebo tests (right)

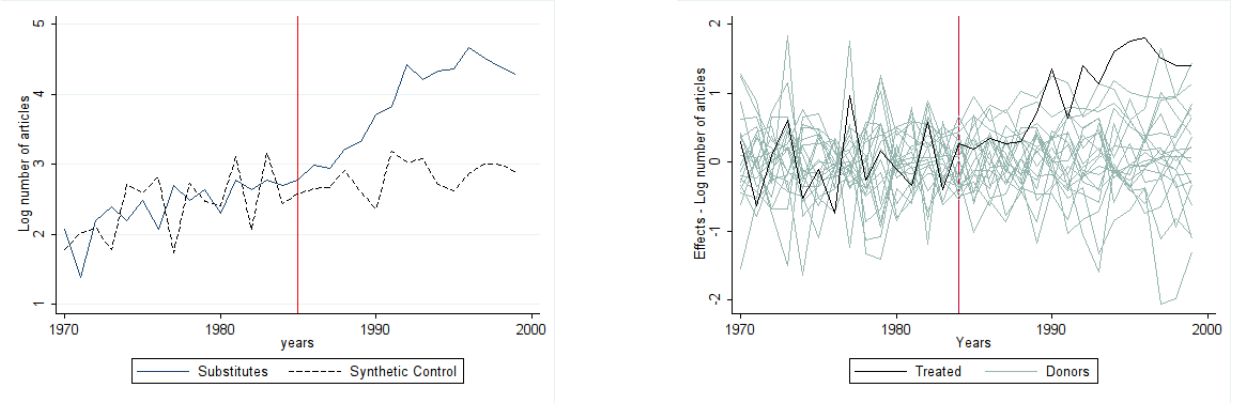

(b) Articles: raw effect (left) and placebo tests (right)

Figure 12: SCM graphs for CFC substitutes assuming anticipation

Note: The graphs display the results of the synthetic control method for substitutes for patents and articles assuming anticipation. For these experiments, the treatment year is redefined as 1985 and the synthetic control constructed using data up to 1982. Results are similar to previous SCM experiments. Specifically, there are no take-offs before 1990. The graphs corresponds to SCM implementations that yielded the lowest pre-RMSPE. That is, for both patents and articles, the SCM uses log count and weighted means of topic proportions. The graphs on the left-hand side represent the raw effect, that is the observed time series of the treated group along with the time series of the constructed control. On the right-hand sides are shown the placebo tests, the non-parametric tests to evaluate the significance of the results; black lines show the effect on the treated group relative to the control group, while each gray line is a placebo test performed on an unit drawn from the donor pool.

Table 7: HAPs contributing to the synthetic control

(a) Patents

HAPs	Weight	Keywords
3,3-Dimethoxybenzidine	0.605	Intermediate: dyes, pigments
o-Xylenes	0.151	Solvent: paints, gasoline
Ethyl chloride	0.114	Solvent, refrigerant, topical anesthetic; leaded gasoline
Pentachlorophenol	0.072	Pesticide, Wood preservative
m-Cresol	0.058	Disinfectant; Preservative; Intermediate: herbicide, explosive

(b) Articles

HAPs	Weight	Keywords
Chloroprene	0.545	Intermediate: adhesives, automotive and industrial parts
Caprolactam	0.230	Intermediate: synthetic fibers, plastics, coatings
1,4-Dichlorobenzene	0.131	Fumigant; deodorant; intermediate: insecticides
Captan	0.093	Fungicide; food, cosmetics, pharmaceuticals

Note: The tables describe the HAPs entering the synthetic control for the main SCM specification (small pool, log counts, weighted means, 5 topics). The information displayed in the "Keywords" column was collected from the EPA.

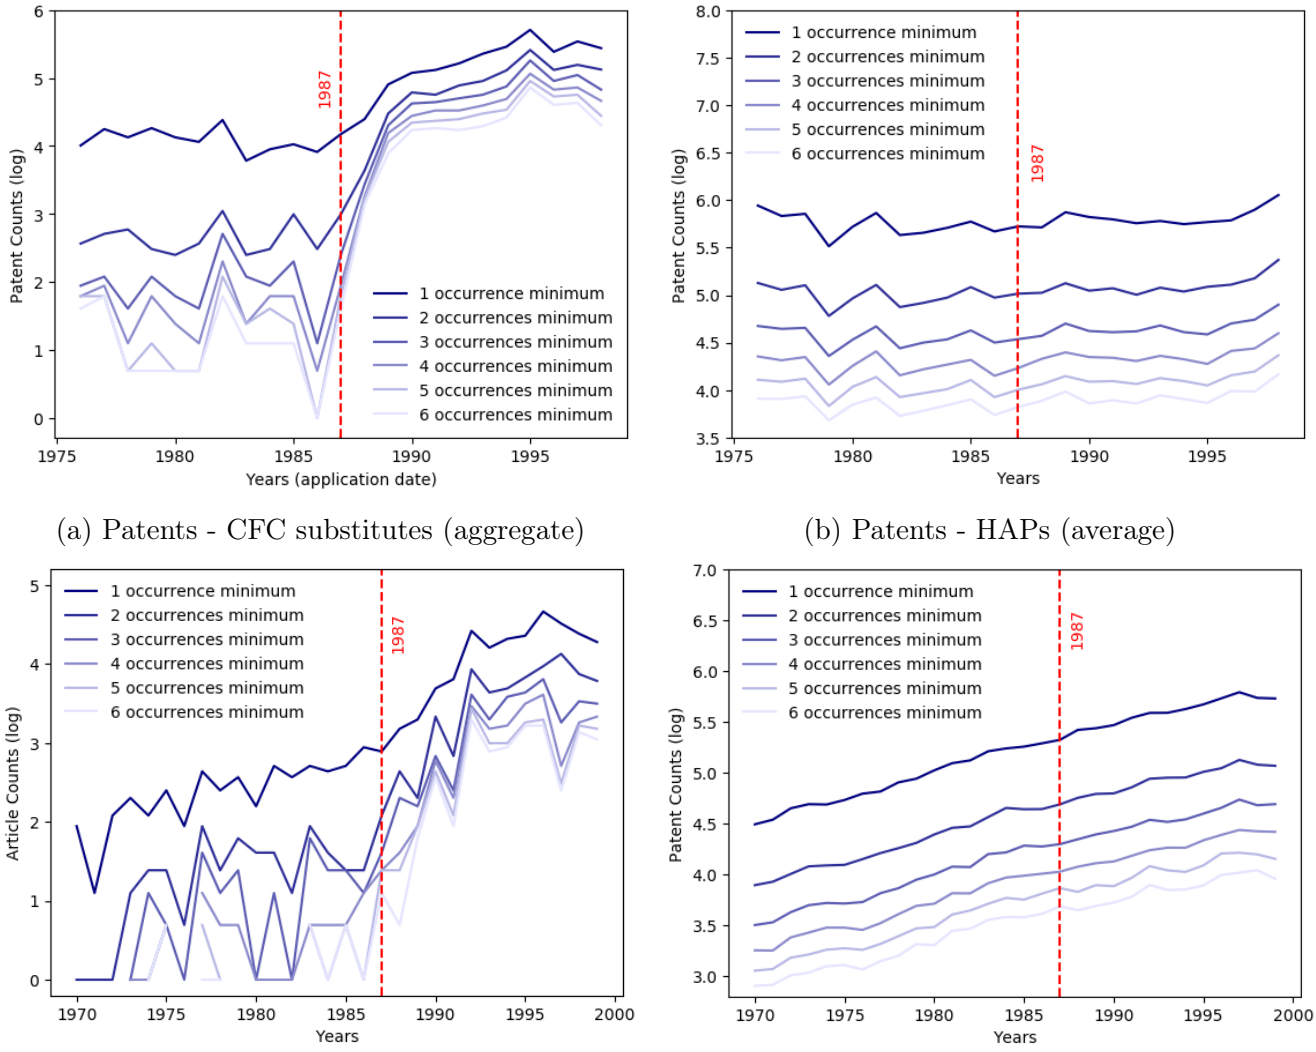

(a) Patents - CFC substitutes (aggregate)

(b) Patents - HAPs (average)

(c) Articles - CFC substitutes (aggregate)

(d) Articles - HAPs (average)

Figure 13: Robustness check: counts with several thresholds of molecule occurrences

Note: The graphs illustrate that focusing on patents and articles that contain more than just one occurrence of molecule would change little to the main analysis. As we increase the occurrence threshold, the trend for the average HAPs remain very similar; only levels decrease. For CFC substitutes, focusing only on patents with greater number of occurrences exarcerbates the differential between the pre and post trends.

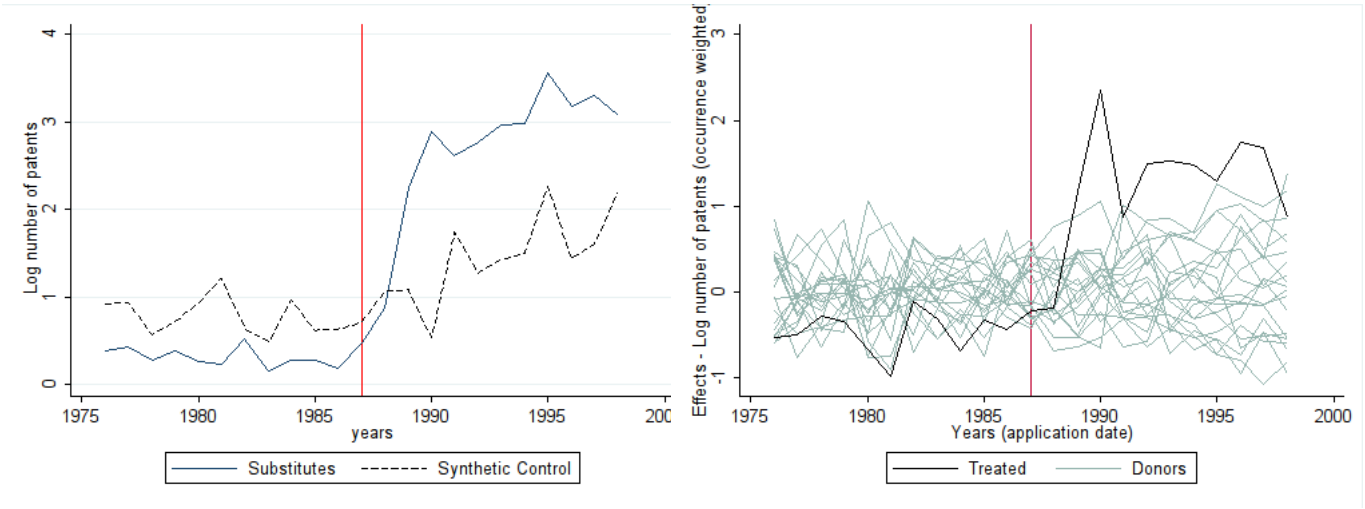

(a) Occurrence weighted

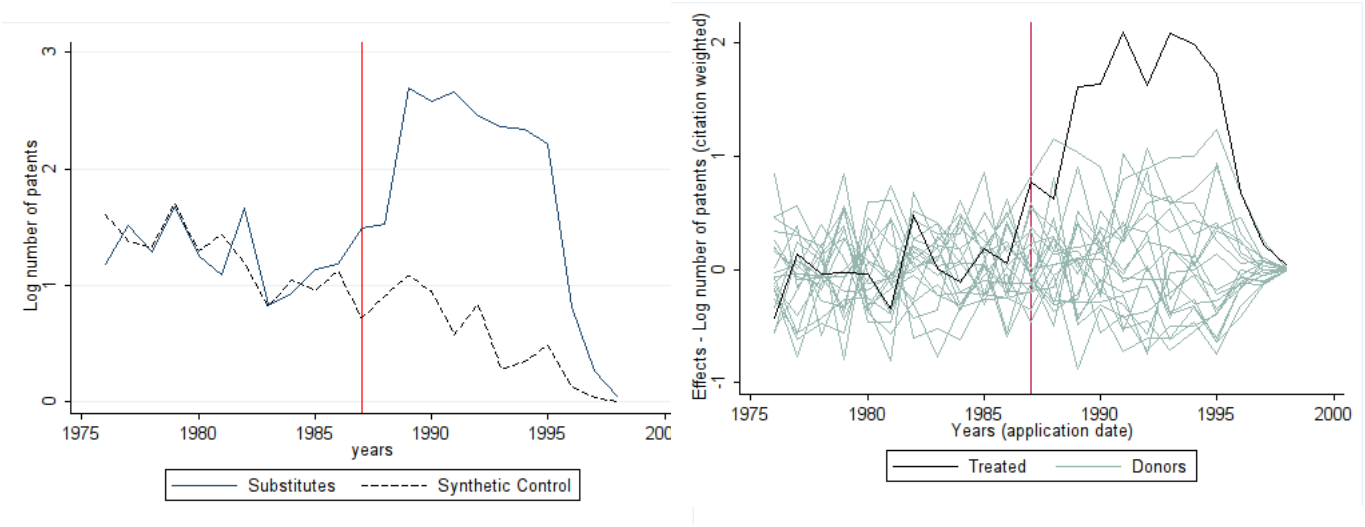

(b) Occurrence and citation weighted

Figure 14: Robustness check: SCM with counts weighted by occurrences and citations

Note: These figures show that implementing the SCM using patent counts weighted by molecule occurences and patent citation does not alter the main conclusions.

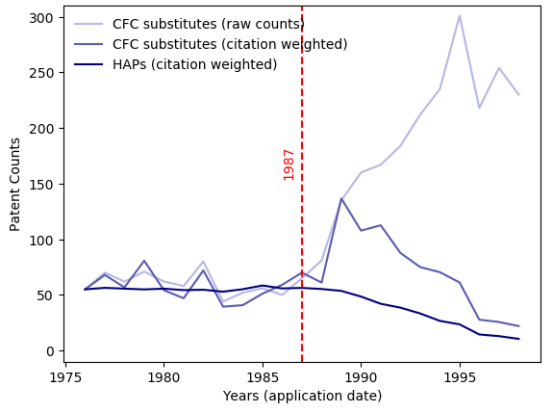

(a) Patents - one occurrence

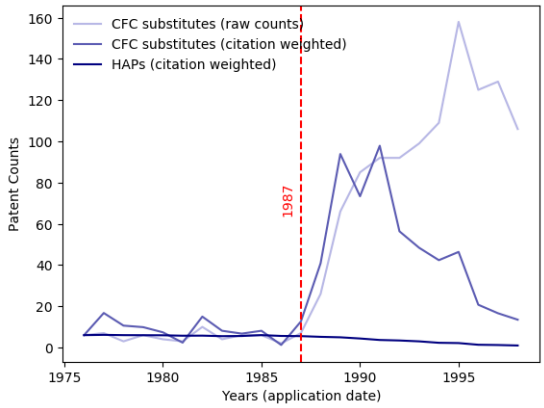

(b) Patents - three occurrences

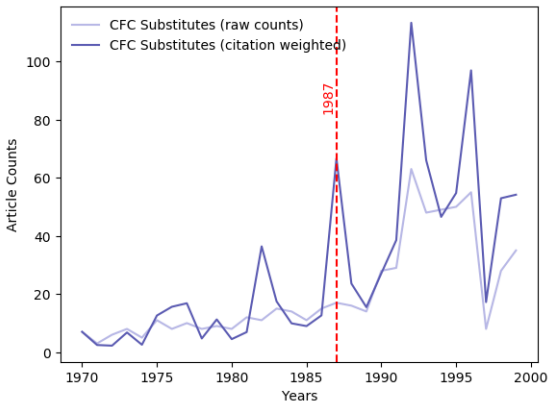

(c) Articles - one occurrence

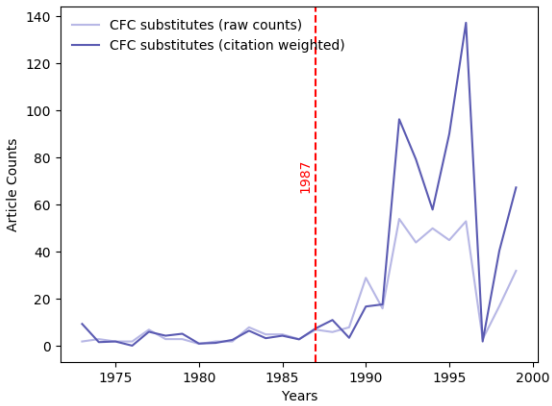

(d) Articles - three occurrences

Figure 15: Time series of citation weighted counts

Note: The graphs illustrate that the most cited articles and patents were published after 1987. The graphs on the left-hand side include any document that mention at least one occurence of a molecule. To test the robustness of this findings, I plot similar graphs but for patents and articles that mention at least three occurences on the figures on the right-hand side. I find that highly cited patents and articles are even more so concentrated after 1987. Collecting citation data for HAPs is undergoing and limited by quotas on the Elsevier API. Additionaly, patent citations include only citations as of 2000; I will be adding citations until 2017.

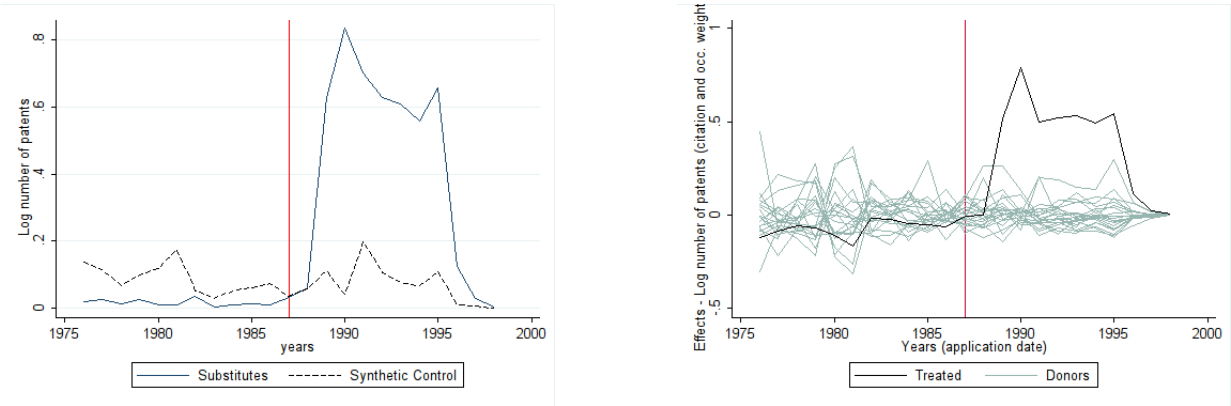

Figure 16: SCM graphs using citation weighted patent counts

Note: The graphs illustrate the robustness of the main results using citation weighted patent counts.

Table 8: Titles of the 10 most cited articles mentioning CFC substitutes

Title	Year	Cited By
Methods for the synthesis of gem-difluoromethylene compounds	1996	333
A new, efficient and environmentally benign system for car air-conditioning	1993	255
High-pressure fluid-phase equilibria: Experimental methods and systems investigated (1988-1993)	1995	227
Evaporation heat transfer and pressure drop of refrigerant R-134a in a small pipe	1998	211
Gas and vapor transport properties of amorphous perfluorinated copolymer membranes based on 2,2-bistrifluoromethyl-4,5-difluoro-1,3-dioxole/tetrafluoroethylene	1996	184
Boiling of new refrigerants: A state-of-the-art review	1996	144
Condensation heat transfer and pressure drop of refrigerant R-134a in a plate heat exchanger	1999	142
Thermochemical and chemical kinetic data for fluorinated hydrocarbons	1995	130
Supercritical fluid extraction in environmental analysis	1993	121
A kinetic study of the reaction of chlorine atoms with CF3CHCl2, CF3CH2F, CFCl2CH3, CF2ClCH3, CHF2CH3, CH3D, CH2D2, CHD3, CD4, and CD3Cl at 295Âś2 K	1992	113

Note: The table displays the titles of the most cited articles mentioning CFC substitutes. Only articles with three molecule occurrences in the text were kept in the sample. We note that these articles, as expected, seem to focus on chemical and physical characteristics of CFC substitutes ("boiling", "evaporation", "pressure" etc...).

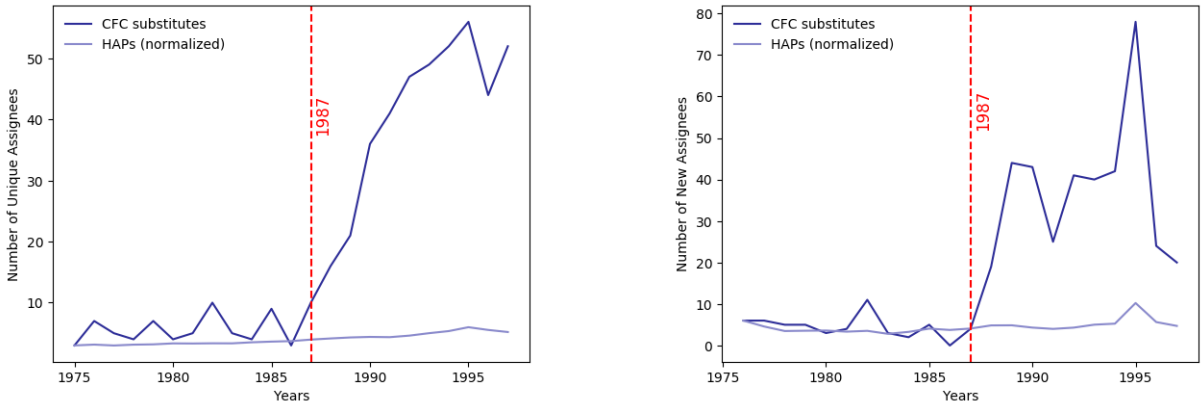

(a) Unique assignees

(b) New entrants

Figure 17: Number of unique and "new entrant" patent assignees

Note: Figure 17a displays the yearly number of unique assignees with patents mentioning CFC substitutes and HAPs. It indicates that the likely presence of new entrants in the post 1987 period. Figure 17b displays the yearly number of assignees that are "new", meaning it is the first time they appear in the data with a patent mentioning CFC substitutes and HAPs. The figure confirms that, after 1987, many firms with no prior experience on CFC substitutes begin patenting. In both figures, only patents with at least 3 molecule occurrences are kept in the sample.

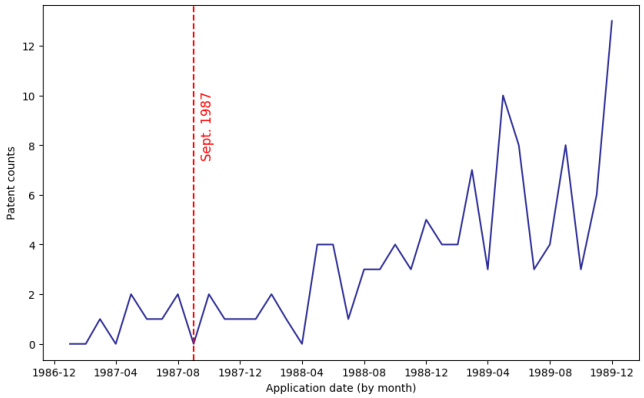

(a) All assignees

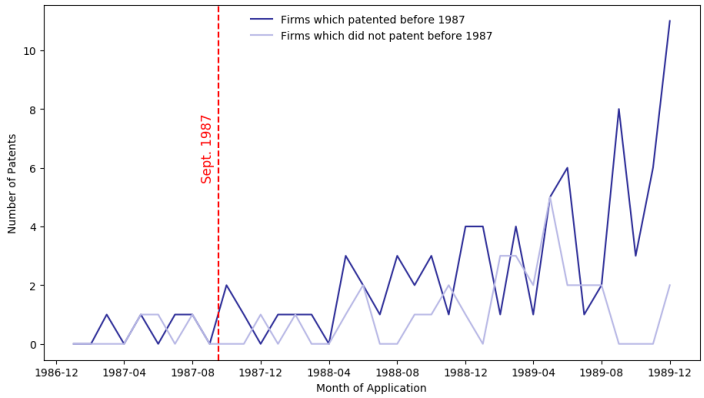

(b) Old vs. new entrants

Figure 18: Monthly counts for patents mentioning CFC substitutes

Note: The graphs show the monthly trends in count of patents mentioning CFC substitutes. It is possible that some firms started working on CFC substitutes before Montreal without seeking patent protection. If those firms has achieved significant advnaces in developping CFC substitutes, we would expect a one-time increase in patent counts in the immediate aftermaths of Montreal. We see on the first graph that this is not the case. Furthermore, if the extent of R&D efforts provided before Montreal was the key driver to the post-Montreal increase in patenting, we should observe major differences in the patenting trends of old and new entrants. On the second graph, I present trends for assignees that never obtained any patent mentioning CFC substitutes before 1987 and those who did. Although a gap seem to build up over time, trends look mostly similar. Only patents with at least 3 molecule occurrences are kept in the sample. The year used is the application year. The period "Before 1987" includes the year 1987.

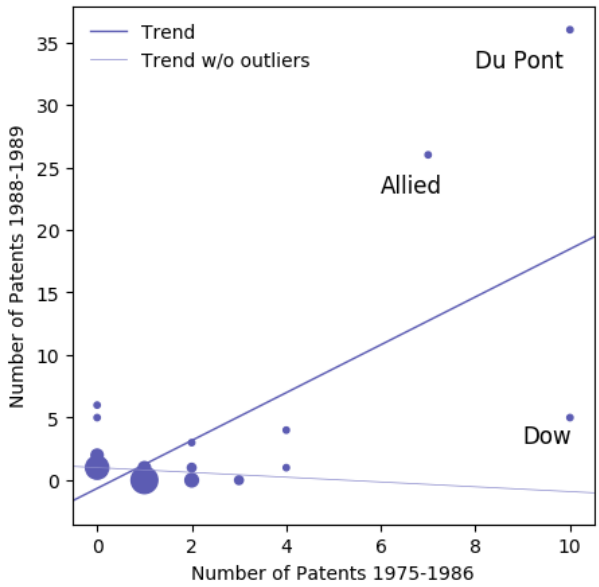

Figure 19: Scatterplot of patenting activity before and after 1987

Note: The figure shows that, on average, firms with more patents before 1986 tend to also have more patents in the immediat aftermaths of Montreal (1988 and 1989). But, the graph illustrates that this effect is largely driven by three outliers: Du Pont, Allied and Dow. Excluding these three firms, there is no clear correlations between patenting prior to 1987 and patenting in the immediat aftermaths of Montreal. The size of the dot is proportional to the number of firms. Only patents with at least 3 molecule occurrences are kept in the sample. The year used is the application year.

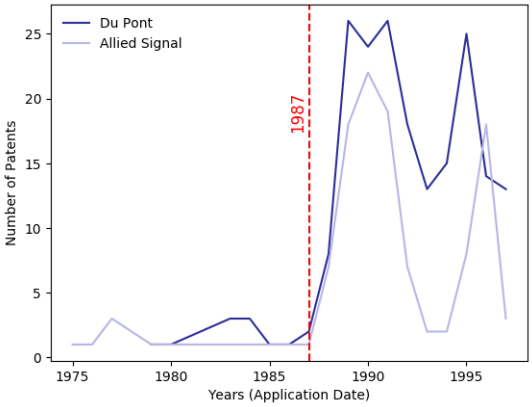

(a) Yearly Counts

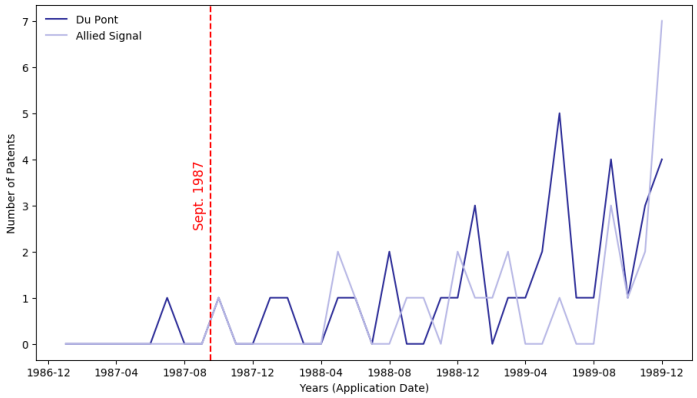

(b) Monthly Counts

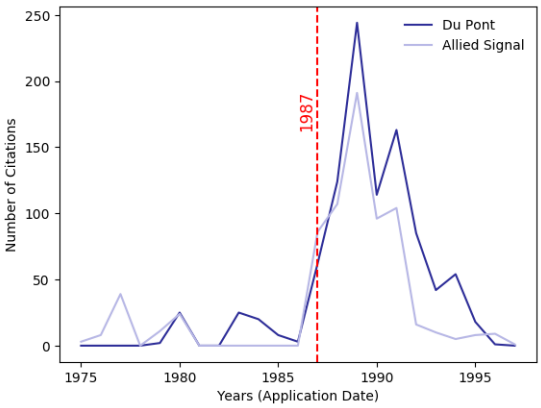

(c) Yearly Citation Weighted Counts

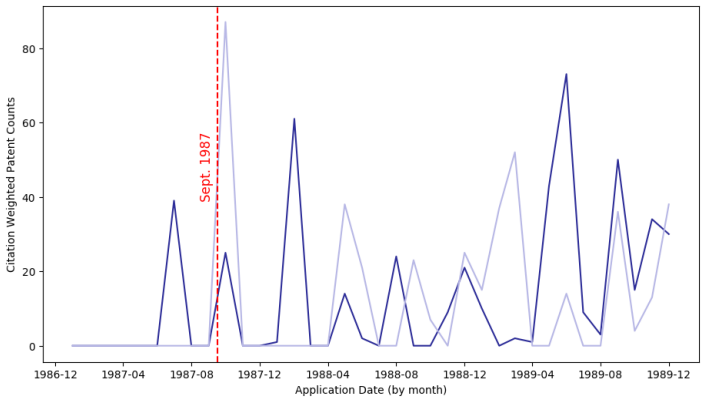

(d) Monthly Citation Weighted Counts

Figure 20: Patent counts for Du Pont and Allied

Note: Figure 20a shows that most patents granted to Du Pont and Allied were applied for after 1989. Figure 20b shows that there is no sudden peak patenting right after Montreal. Instead we observe a gradual ramping up of patenting activity. Figure 20c illustrates that the patents granted to Du Pont and Allied which received the greatest number of citations mostly originate from 1989 to 1991. Figure 20d indicates, however, that, in the weeks that followed Montreal, both Du Pont and Allied applied for patents that would go on receiving a high number of citations. Only patents with at least 3 occurrences of a molecule are retained in the sample.

A Appendix A

A1 Cleaning procedure

A1.1 Full-text

For most articles, the full text downloaded from ScienceDirect requires many different cleaning steps before it can be analyzed. The text is often the imperfect result of the conversion of images of typed or printed text into machine-encoded text: some words are not well recognized especially when the article contained mathematical symbols and equations. Words are also sometimes not properly separated by space. Below are the successive cleaning steps I undertake. Patent texts require only a few preprocessing steps described below under the *preprocessing* section.

Preprocessing. I fix "broken" unicode such as garbled HTML entities, convert non-ascii characters into their closest ascii equivalents, replace all URL strings with "URL", replace all email strings with "EMAIL", replace all phone number strings with "PHONE", replace all currency symbols with their standard 3-letter abbreviations, replace English contractions with their non-shortened forms, replace all accented characters with unaccented versions. This is done through the use of the Python package *Textacy*. I also remove any sequence of more than 1000 digits. I also remove tokens with more than 50% of characters being digits.

Drop non-English articles. Some articles seem not to be written in English. For this reason, I use Google's CLD2 library in Python to detect every document's language, and drop those that are detected with large enough confidence as not being English.

Splitting text into sentences. I use the Python package *Spacy* to parse the text and detect sentences.

Clean each sentence. I remove any punctuation and tokens of length 1 (i.e. stand-alone characters). I also drop sentences if their number of non-digit tokens is lower than 5. This helps remove unintelligible sequences of letters and digits that are often found at the beginning of articles. Finally, I remove the entire sentence if less than 80% of tokens are recognized by SpaCy's English dictionary. This provides a rough test for whether the sentence is written in another language. Indeed articles can sometimes present translations in other languages within the full text.

Further processing. Number-like strings are replaced with the token "*NUMBER*". All words are lowercased.

A1.2 Meta-Data

Scopus's meta-data provides the name and geographic localization of the authors' affiliations. However, Scopus does not provide information about these organization. In particular, knowing the share of articles affiliated with public vs. private entities would be interesting.

To that aim, I leverage the Global Research Identifier Database[68] (GRID) which provides information about a worldwide collection of organizations associated with academic research. In

[68] `https://www.grid.ac/`

particular, GRID classifies an entity as one of the following types: education, company, government, facility, non-profit, health care[69]. An organization is classified as "education" if it can grant degrees, as "company" if it is a business entity with the aim of gaining profit, as "government" if it is operated mainly by a government, and as "health care" if it is a place that treats patients. Facilities encompass building or facilities researching specific areas and usually containing specific equipment (e.g., a nuclear plant). Nonprofits include charities but also non-governmental research institutes[70].

Unfortunately, the name of the organizations and its geographical location are often reported differently in Scopus and GRID. To match as many entities as possible, I first look for exact matches, then for approximate ones using tools such as fuzzy matching in python. Still, many remained unmatched. I then manually match any organization appearing, at least, three times or more in the data. There were about 300 of such organizations.

The bulk UPSTO downloadable data contains patent meta-data. Names and addresses of the inventors and assignee are therefore more readily available. I use the country of the assignee, and when the patent has no assignee, I use the country of the inventor. The USPTO data, however, does not classify assignee concerning the type of the organization (e.g., company, education or non-profit). The GRID database here is not as useful because most patents originate from businesses; GRID encompasses some for-profit entities with major research activities, but many patentees are in fact small companies unlikely to be listed under GRID. Hence, to match patent assignees to a type, I implement a more basic strategy. It is useful to notice that the name of an organization tends to contain tokens informing about the nature of that organization. For example, the "Inc." abbreviation in the name *Flow Vision, Inc.* tells us that this is a for-profit organization. Other such tokens includes "corp.", "co.", "plc", "llc", "limited" or "company", as well as "& cie"[71]. Similarly, I identify organizations containing tokens such as "university" or "school" as being of the "education" type, and those containing tokens such as "govern", "ministr" or "agency" as being of the "government" type. The use of these simple rules helps me match about 36529 out of 45820 assignee names. Out of the 7899 remaining, I manually match those that appear at least ten times in my data (about 200 of them). I leave the rest with no type information.

A2 Topic Modeling

I use topic modeling, a machine learning method for text analysis, to generate covariates that describe the semantics surrounding molecules and therefore proxying some chemically and industrial characteristics. Specifically, I use Latent Dirichlet Allocation (LDA), a method of probabilistic topic modeling for text (Blei 2012; Blei et al. 2006, 2009; Roberts et al. 2014; Roberts et al. 2016).

In this method, the experimenter chooses the number of topics, and after training the algorithm on a corpus, the model can return the topic distribution for each document. Put differently, using the words that appear in a given document; the LDA model infers what proportion of each topic a document contains. Intuitively, the topic proportions describe quantitatively what an article talks about, and we can, therefore, think of it as a proxy of the physical, chemical and industrial characteristics of a molecule. I train the algorithm, not on the entire corpus, but on the subset

[69]There are two other classifications: "archive" and "other." For more information, see `https://www.grid.ac/pages/policies`

[70]For example, in the USA, the National Academy of Sciences is classified as a non-profit.

[71]In other languages, here are a few of the tokens that I found in the data: "kaisha" or "kk" in Japanese, "spa" in Italian, "gesellschaft" or "gmbh" or "ag" or "kg" in German, "bv" or "nv" in Dutch, "sa" or "sarl" in French, "ab" in Swedish, "oy" in Finnish, "rt" in Hungarian.

of documents that contain at least one mention of a molecule: this represents a total of 382,599 patents and 382,005 articles. The LDA model is trained choosing five topics. Table A1 displays the top three words in the five topics generated by the LDA model on the corpus of patents. The supporting online material features the lists of top 20 words in each topic for patents and articles.

I then aggregate the topic proportions of the documents at the molecule level by calculating a weighted mean topic proportion with weights proportional to the number of times an article mentions a molecule. As a result, articles with many mentions of a molecule contribute more to the aggregated topic proportion. I also test the robustness of my results to taking a simple non-weighted mean. Figure 2 summarizes these various steps with a simple example of three documents, two molecules, and two topics.

Finally, I use these topic proportions together with the outcome variable (log count) as covariates in the synthetic control method. Hence, the algorithm will construct a synthetic control that not only reproduces the path of log count in pretreatment periods but that also mimics the values of the different topic proportions. Figure A1 displays the correlation heat map between the topic proportions and counts in patents and articles. We note that the topic proportions are only weak predictors of the variable count. Hence, in the SCM optimization, they should have a small contribution in constructing the synthetic control because the SCM algorithm assigns bigger weights to variables that are better predictors. Figure 9, however, illustrates why topic proportions are still useful. The graphs display scatter plots of topic proportions and log count. We see that some HAPs have values of topic proportions that stand out as outliers. This indicates that those HAPs present a semantic context that is likely very different from the one of CFC substitutes. Hence topic proportions ensure that such HAPs are not used in constructing a synthetic control.

Table A1: Top 3 words in topics for patents

	Words	Probability
Topic 1	crotononitrile	0.0090
	remote	0.0063
	dialkylhydantoin	0.0047
Topic 2	andreu	0.0141
	sulfon	0.0075
	phosphatidylinositols	0.0072
Topic 3	neal	0.0323
	isopropyltrimethoxysilane	0.0276
	inducers	0.0236
Topic 4	topcoatings	0.0071
	heterophasic	0.0054
	neal	0.0052
Topic 5	trisethyl	0.0157
	maker	0.0128
	amineprotecting	0.0066

Note: The table presents the three most probably words in the five topics generated by the LDA algorithm. As a consequence of the nature of the corpora (patents), the words in the topics are highly technical and specialized which makes it difficult to associate one topic to a general theme.

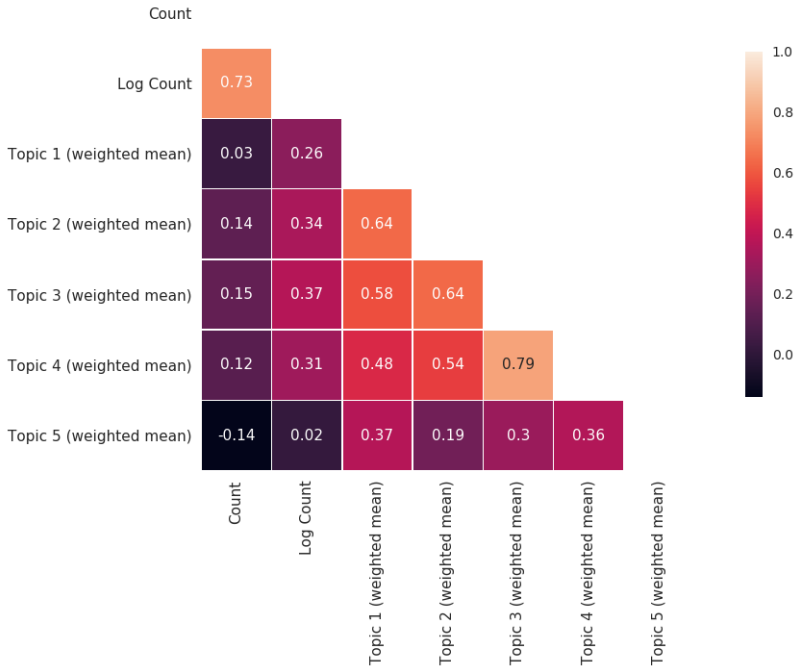

(a) Patents.

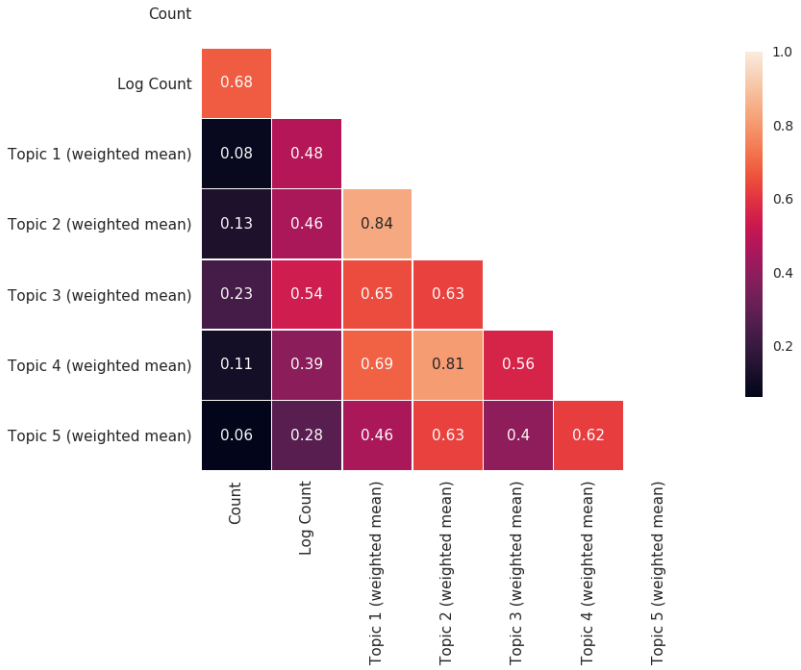

(b) Articles.

Figure A1: Correlations between topic proportions and counts

Note: We see that topics somehow correlate with log counts. The heatmaps used the whole sample of HAPs.

A3 SCM theoretical foundations

Here, I briefly summarize the theoretical underpinnings of the SCM. Suppose there are J+1 molecules, J molecules as potential controls and one, denoted with the subscript 1, that is treated. The treatment effect can be written as $\alpha_{it} = Y_{it}^T - Y_{it}^N$, where Y_{it}^N is the number of document mentioning molecule i in year t if no intervention, and Y_{it}^T the number of documents mentioning molecule i in year t if intervention. Here the quantity we need to estimate is Y_{it}^N. Abadie et al. (2010) show that a weighted average of the control units can approximate the counterfactual Y_{it}^N, that is:

$$Y_{1,t}^N \rightarrow \sum_{j=2}^{J+1} w_j^* Y_{jt} \text{ with } w* \text{ s.t. } \sum_{j=2}^{J+1} w_j^* Y_{jt} = Y_{1,t} \text{ and } \sum w_j^* Z_j = Z_1$$

To understand why this is the case, Equation 5 presents the underlying factor model. δ_t is an unknown common factor w constant loadings across units; θ_t is a vector of unknown parameters; Z_i a vector of observed covariates (not affected by intervention); λ_t unobserved common factors; μ_i a vector of unknown factor loadings and ϵ_{it} unobserved transitory shocks with zero mean. Note that this model generalizes the difference-in-differences model which imposes that λ_t be constant for all t. Hence, the unobserved confounders are constant in time and can be eliminated by taking time difference. Here, the SCM allows the effects of confounding unobserved characteristics to vary with time; taking time differences would not get us rid of μ_i.

$$Y_{it}^N = \delta_t + \theta_t Z_i + \lambda_t \mu_i + \epsilon_{it} \tag{5}$$

A synthetic control such that $\sum_{j=2}^{J+1} w_j^* Z_j = Z_1$ and $\sum w_j^* \mu_j = \mu_1$ would be unbiased estimator of Y_{1t}^N. In other words, fitting Z_1 and Y_{11} ... Y_{1T_0} is a way of indirectly fitting μ_1, the unobserved factor loadings. As a result, it is important to restrict the donor pool to units with outcomes that are thought to be driven by the same structural process as for unit representing the case of interest and that were not subject to structural shocks to the outcome variable during the sample period.

B Appendix B: Additional Figures and Tables

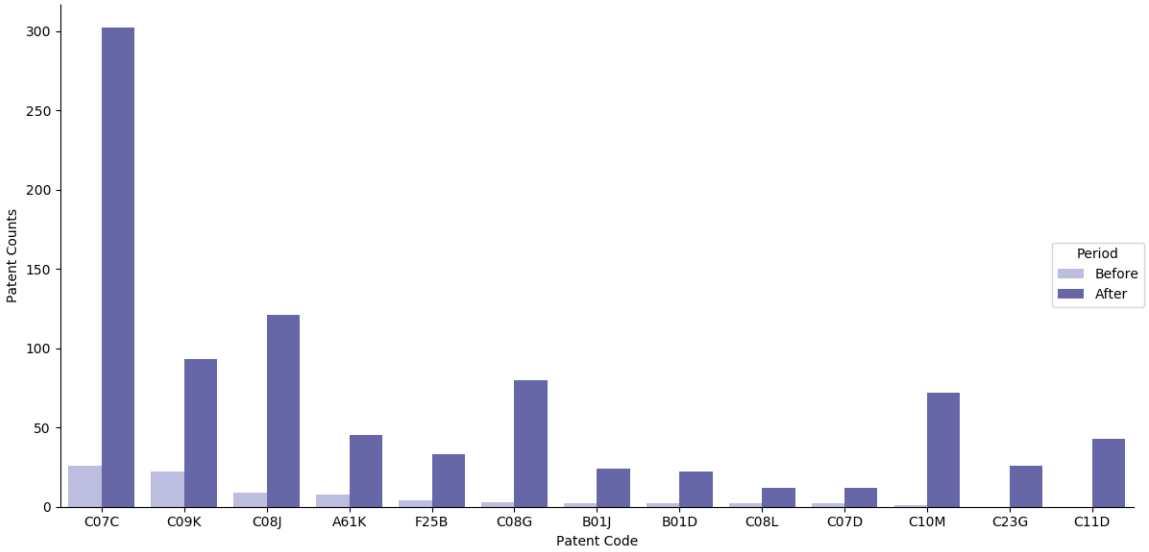

Figure B1: Most frequent codes for patents mentioning CFC susbtitutes before and after 1987

Note: The figure illustrates the differences between the most frequent codes for patents before and after 1987 (year of application is used). The most frequent patent codes before 1987 tend to be the most frequent after 1987. At the same time, some codes with low to zero frequency before 1987 become important after 1987 (e.g., C08G, C10M, C23G or C11D). Only patents with at least 3 molecule occurrences are kept in the sample.

Table B1: List of CFC substitutes

Substitute	PAFT	AFEAS	Substitute for
HCFC-22	No, already marketed, toxicology known	Yes	CFC-11, CFC-12 in foams
HCFC-142b	No, already marketed, toxicology known	Yes	CFC-11, CFC-12 but not ideal
HFC-152a	No, already marketed, toxicology known	Yes	CFC-11, CFC-12 but not ideal
HCFC-123	Yes	Yes	CFC-11 in refrigeration
HFC-134a	Yes	Yes	CFC-12 in refrigeration (car AC)
HCFC-141b	Yes	Yes	CFC-11 in foams
HCFC-124	Yes	Yes	CFC-114 in refrigeration and sterilization
HCFC-125	Yes	Yes	CFC-115 in refrigeration and sterilization
HCFC-225ca	No, second rank candidate	Yes	
HCFC-225cb	No, second rank candidate	Yes	
HFC-32	No, second rank candidate	Yes	refrigeration
HFC-143a	No, second rank candidate	Yes	CFC-12 in refrigeration
HFC-245fa	No	No	CFC-11, HCFC-141b and HCFC-142b in foams
HFC-365mfc	No	No	CFC-11, HCFC-141b and HCFC-142b in foams

Note: The table lists 14 molecules that were considered as potential CFC substitutes in 1988. The columns PAFT and AFEAS indicate whether the molecule was included in the investigations carried out by the PAFT and AFEAS. The PAFT (Program for Alternative Fluorocarbon Toxicity Testing) was created in January 1988 to work on assessing the toxicity of five possible alternatives. The AFEAS (Alternative Fluorocarbon Environmental Acceptability Study), created in December 1988, investigated the atmospheric dynamics of twelve potential CFC substitutes. I use these twelve molecules to form the group of CFC substitutes. I also include in this group two other possible CFC substitutes mentioned in Benedick (2009) and Parson (2003). In the rest of my analysis, I track the evolution of patents and articles mentioning these 14 molecules.

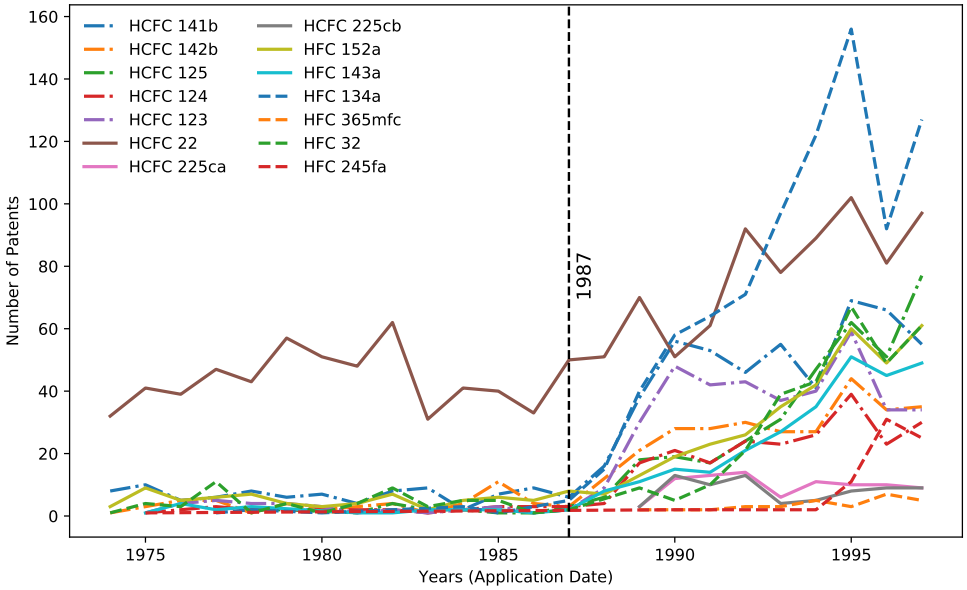

(a) Patents

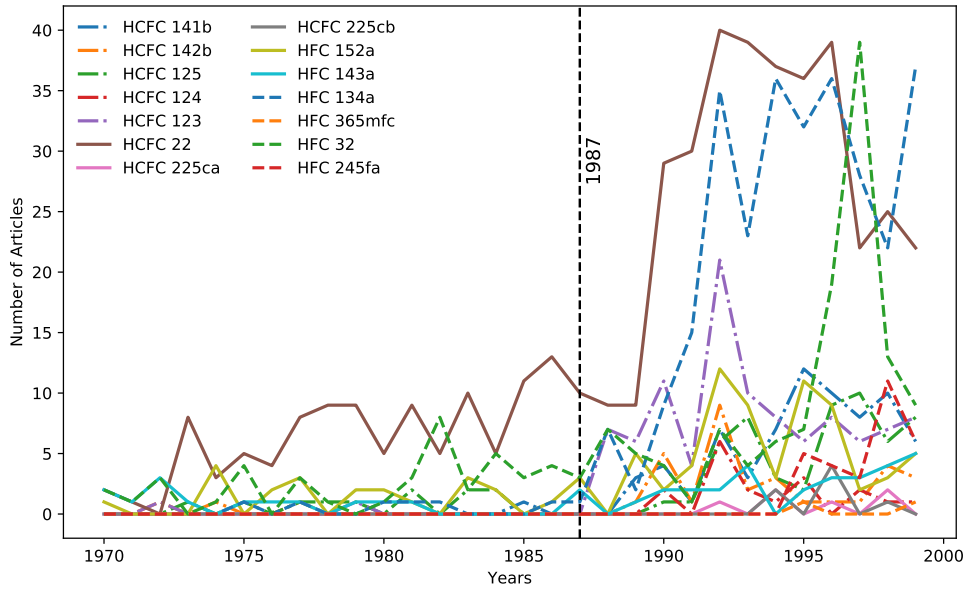

(b) Articles

Figure B2: Counts in articles and patents for each CFC substitute

Note: These graphs plot the yearly number of articles or patents mentioning the names of given CFC substitutes. We note a clear increase for most CFC substitutes in the 1990s. For patents, the graph shows patents that have been *granted* (as opposed to patent applications) but the years on the x-axis corresponds to the application date. There is on average a two-year delay between patent application and grant.

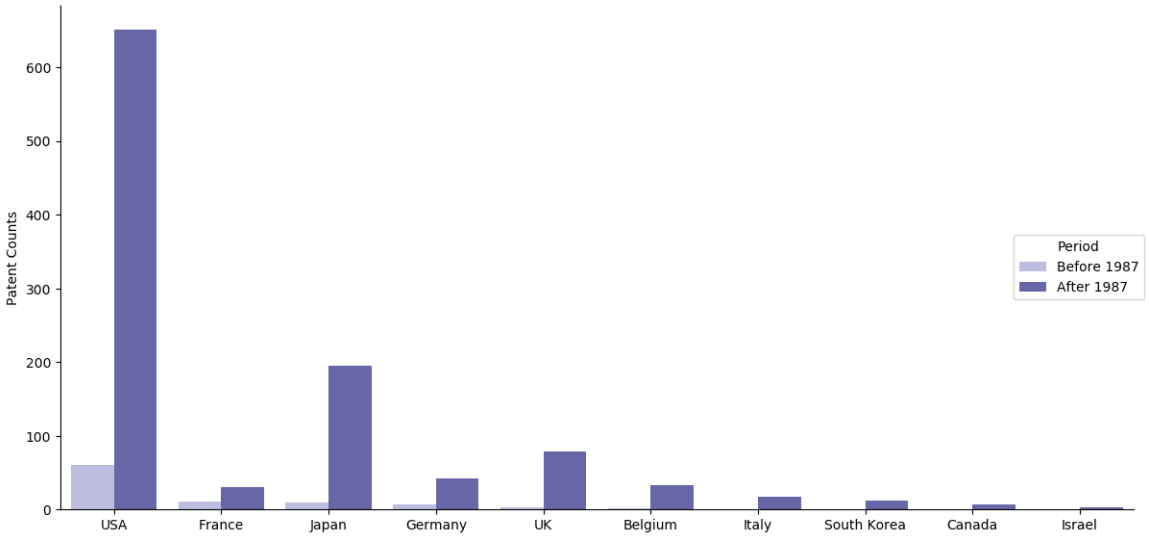

Figure B3: Country of origin of patent assignees before and after 1987.

Note: The figure illustrates that all countries are associated with an increase number of patents mentioning CFC substitutes. We note in particular the strong increase for Japan and the UK. Only patents with at least 3 molecule occurrences are kept in the sample. The year used is the application year. The period "Before 1987" includes the year 1987.

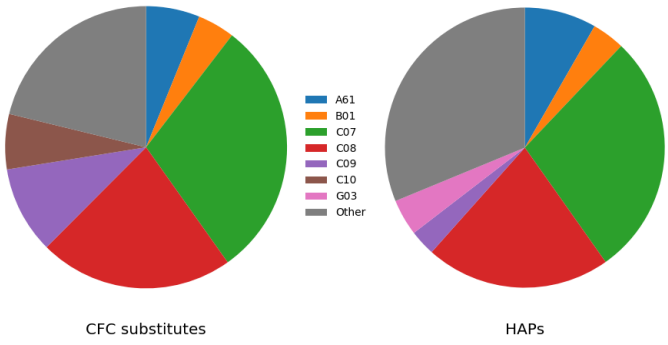

Figure B4: Second-level patent codes

Note: We see that CFC substitutes and HAPs also share similar second-level patent codes.

Table B2: Summary statistics of meta-data

(a) Patents

	CFC Substitutes	HAPs
Education	0.02	0.02
	(0.14)	(0.14)
Company	0.97	0.96
	(0.17)	(0.20)
Government	0.01	0.02
	(0.07)	(0.13)
Facility	0.00	0.00
	(0.07)	(0.02)
Nonprofit	0.00	0.00
	(0.00)	(0.07)
Healthcare	0.00	0.00
	(0.00)	(0.02)
USA	0.61	0.56
	(0.49)	(0.50)
Europe	0.21	0.23
	(0.40)	(0.42)
Japan	0.17	0.17
	(0.37)	(0.38)

(b) Articles

	CFC Substitutes
Citation Count	30.60
	(72.19)
Number of Authors	2.95
	(2.85)
Education	0.75
	(0.43)
Company	0.12
	(0.32)
Government	0.11
	(0.31)
Facility	0.11
	(0.31)
Nonprofit	0.02
	(0.13)
Healthcare	0.02
	(0.15)
USA	0.36
	(0.48)
Europe	0.39
	(0.49)
Japan	0.12
	(0.32)

Note: The tables display summary statistics on the country of origin and types of patent assignees (left-hand side) and on the country and type of affiliation of the authors of articles (right-hand side). Data collection for HAPs in articles is still undergoing due to quota limitation on the Elsevier API. We note that more than 96% of patents are granted to for-profit organizations. The rest is shared among organizations coming from the educational and governmental sector as well as organizations that fit the description of "facility". The majority of patents are granted to assignee domiciliated in the United States. European assignees tend to represent around 20 to 30% of patents; Japanese around 10 to 20%.

For patents, the variables *Education*, *Company*, *Government*, *Facility*, *Nonprofit*, and *Healthcare* are binary variables identifying the type of patent assignee. For articles, "Education" is a dummy variable that equals 1 if at least one of the authors is affiliated with an organization in the higher education sector. "Company" is a dummy variable that equals 1 if at least one of the authors is affiliated with a for-profit private organization. "Government" is a dummy variable that equals 1 if at least one of the authors is affiliated with a governmental entity. "Facility" is a dummy variable that equals 1 if at least one of the authors is affiliated with a facility pursuing research in specialized areas (e.g. nuclear plant, particle accelerators etc...). "Nonprofit" is a dummy variable that equals 1 if at least one of the authors is affiliated with a nonprofit research institute. "Healthcare" is a dummy variable that equals 1 if at least one of the authors is affiliated with an organization where patients are treated. "USA", "Europe" and "Japan" are dummy variables that equal 1 if at least one of the authors is affiliated with, respectively, the USA, a European country and Japan. By European country, I mean any country belonging to the EU in 2016 plus Switzerland, Norway, Serbia, Ukraine, Moldova and Russia. Organization types were collected from the Global Research Identifier Database.

Table B3: Summary statistics for CFC substitutes and HAPs

Variables (pre-1986 average)	Substitutes	HAPs Mean	HAPs Min	HAPs Max	HAPs Std.Dev.
Count	64.2	71.37	41.8	101.8	21.23
Topic 1 (weighted mean)	0.2	0.2	0.16	0.23	0.02
Topic 2 (weighted mean)	0.1	0.11	0.08	0.12	0.01
Topic 3 (weighted mean)	0.43	0.43	0.39	0.47	0.02
Topic 4 (weighted mean)	0.18	0.19	0.15	0.23	0.02
Topic 5 (weighted mean)	0.09	0.08	0.06	0.12	0.02
Topic 1 (unweighted mean)	0.2	0.2	0.17	0.22	0.01
Topic 2 (unweighted mean)	0.1	0.11	0.1	0.12	0.01
Topic 3 (unweighted mean)	0.43	0.43	0.41	0.46	0.01
Topic 4 (unweighted mean)	0.19	0.19	0.17	0.22	0.01
Topic 5 (unweighted mean)	0.09	0.08	0.06	0.1	0.01

Note: The table displays summary statistics for the aggregated CFC substitutes and HAPs for patents. The SCM imposes that the synthetic control's weights be non-negative and sum to 1. This can be problematic in cases where the treated unit lies at the extremes of the distribution of the donor units. We see, here, that the values for CFC substitutes always fall within the range of the values for HAPs. This table confirms that we are not in this case. Hence, there should be no penalty constraining the weights to non-negative and to sum to 1. Note that only HAPs in the small donor pool are used it. Similar results were obtained for articles.

Table B4: Variable weights used in the construction of the synthetic control

	Variable Weight
Topic 1 (weighted mean)	0.04
Topic 2 (weighted mean)	0.04
Topic 3 (weighted mean)	0.02
Topic 4 (weighted mean)	0.03
Topic 5 (weighted mean)	0.02
Log Count	0.86

Note: The table displays the weights assigned to variables in the optimization procedure of the SCM. These weights are for the case of patents with log counts and weighted means of topic proportions, using the small pool of HAPs as donor pool. We note that topic proportions contribute about 15% in constructing the synthetic control.

Table B5: Means over pre-treatment periods for CFC substitutes

	Real S	Synthetic S	Average HAPs
Count (log)	4.17	4.17	4.22
Topic 1 (weighted mean)	0.20	0.20	0.20
Topic 2 (weighted mean)	0.10	0.10	0.11
Topic 3 (weighted mean)	0.43	0.42	0.43
Topic 4 (weighted mean)	0.18	0.19	0.19
Topic 5 (weighted mean)	0.09	0.09	0.08

Note: The table illustrates how the SCM is able to construct a better comparison unit than simply using the mean of many control units. The table displays the mean over the years 1970 to 1985 for log patent counts and topic proportions for the group of CFC substitutes ("Real S"), for the constructed synthetic substitute ("Synthetic S") and for the average of HAPs. The synthetic control, here, was constructed based on similarity with the variables "Log Count" and the weighted means of the 5 topic proportions. We see that the synthetic control matches the real substitute group much better than the average of HAPs in terms of log count. This is the core idea motivating the use of the SCM. The HAPs used in calculating the average are only those from the small pool, explaining why the topic proportions are very similar.

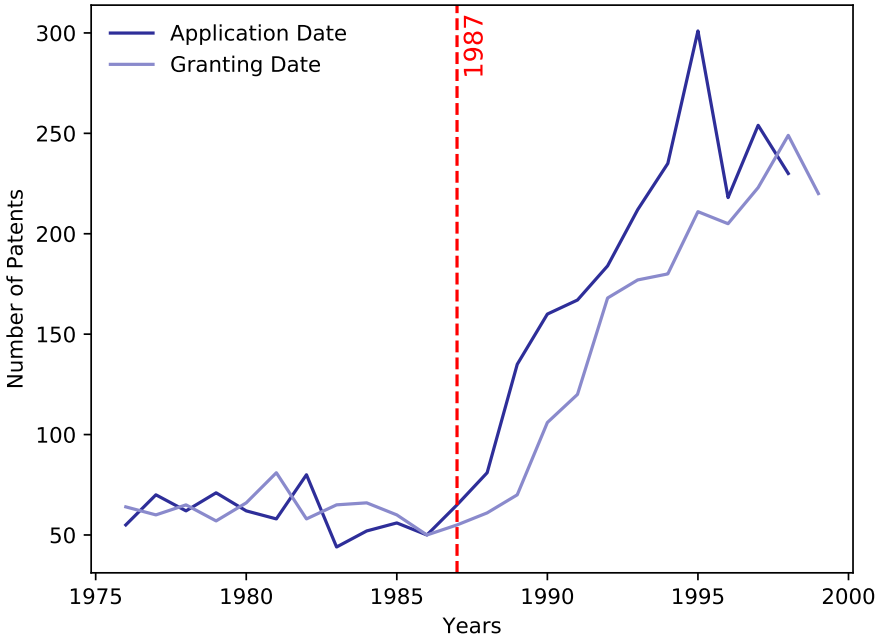

Figure B5: Patent counts with application date vs. grant date for CFC substitutes

Note: The graph plots the number of patents mentioning any CFC subtitutes using the application date of the patent or the granting date. The two curves are very similar, with only about a two-year delay after 1987. Precisely, there is on average a 22-month delay between application and granting, with a standard deviation of 12 months. The graph illustrates that we obtain similar reszults by using the application date for the main analysis.

Table B6: Robustness checks: SCM with counts as outcome variable (instead of log counts)

(a) Patents

Rule	Topic Means	Donor Pool	Pre RMSPE	p-value	ATE	Year
intermediate	unweighted	whole sample	3.35	0.006	71.88	1990
intermediate	unweighted	small pool	3.41	0.000	70.87	1990
intermediate	weighted	small pool	3.70	0.000	72.28	1990
intermediate	weighted	whole sample	3.92	0.000	71.24	1990
weak	weighted	whole sample	8.74	0.000	100.92	1992
weak	unweighted	whole sample	9.68	0.000	84.52	1992
weak	weighted	small pool	10.23	0.000	109.23	1990
weak	unweighted	small pool	10.86	0.000	113.61	1990

(b) Articles

Rule	Topic Means	Donor Pool	Pre RMSPE	p-value	ATE	Year
intermediate	weighted	small pool	1.54	0.000	35.17	1990
intermediate	unweighted	small pool	1.78	0.000	35.52	1990
intermediate	unweighted	whole sample	1.96	0.000	35.04	1992
intermediate	weighted	whole sample	2.03	0.000	35.39	1992
weak	weighted	small pool	2.65	0.000	44.08	1990
weak	unweighted	small pool	2.86	0.000	42.6	1990
weak	weighted	whole sample	4.44	0.006	45.81	1998
weak	unweighted	whole sample	5.65	0.006	46.73	1994

Note: The tables display results for when using counts in levels instead of in log as the outcome variable. The estimated treatment effects for these robustness checks are similar to the effect estimate with the main methodology. "Topic Means" indicates the procedure for aggregating the topic proportions at the molecule level. If "weighted", the calculated proportion of topic j for molecule i is the mean proportion of topic j across all documents mentioning molecule i, weighted by the number of times the molecule appears in the document. "Donor Pool" indicates what sample of HAPs is used in the SCM procedure. For "small pool", the sample of HAPs used corresponds to the twenty HAPs most similar to the treated unit in terms of counts and topic proportions before 1987.

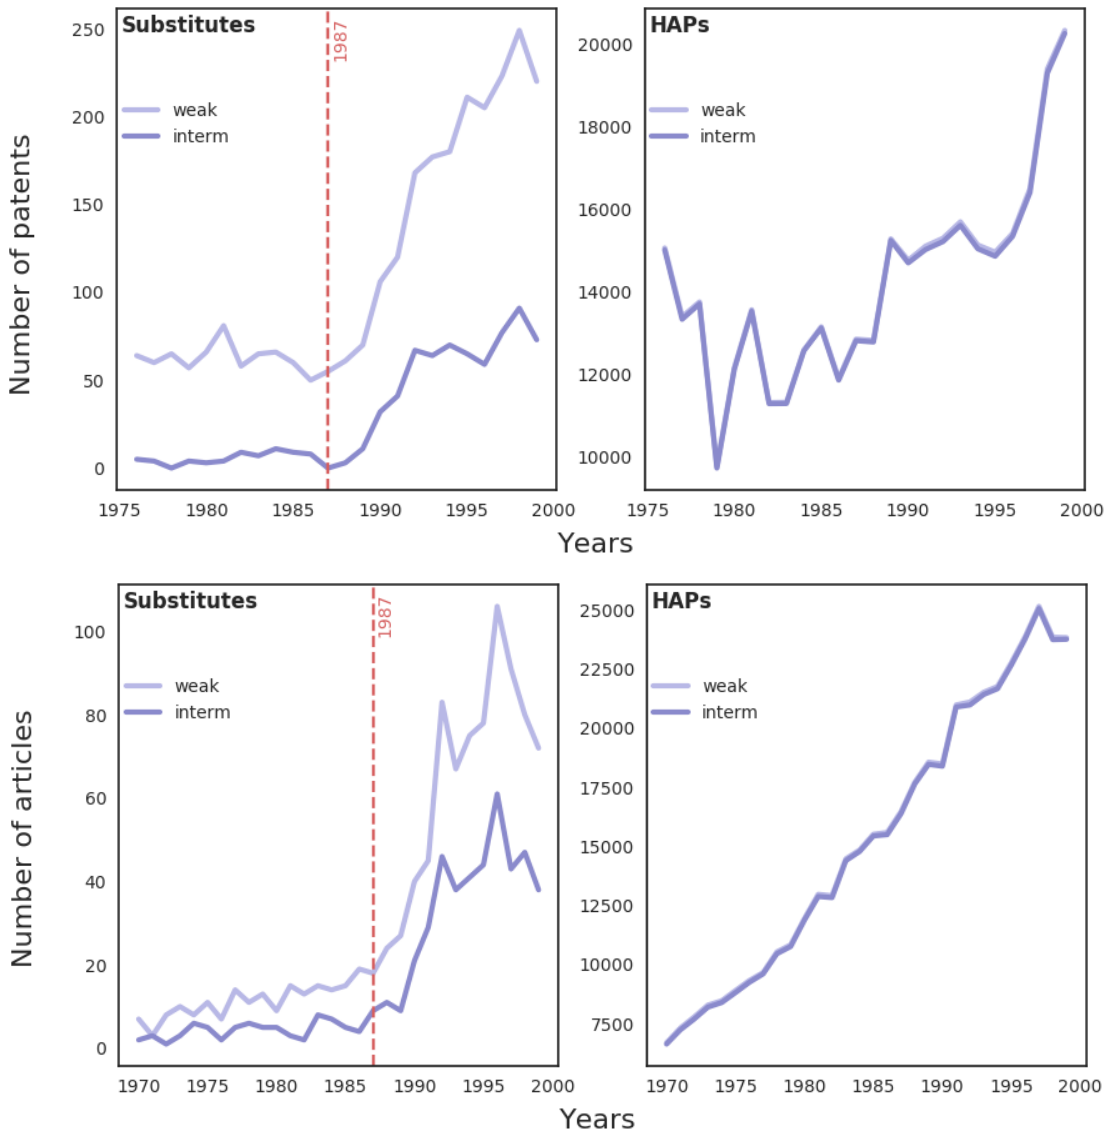

Figure B6: Yearly counts for CFC substitutes and HAPs with different assignment rules for patents and articles.

Note: The figures show the number of patents and articles each year for CFC substitutes and HAPs according to the two different assignment rules. The main results were obtained using the rule called "weak". I test the robustness of these results by using an alternative rule, which I call "intermediate". The intermediate rule is a more conservative way of assigning documents to molecules. When a document mentions several molecules, instead of assigning the document to each molecule, the document is assigned to only the molecule it mentions the most. We note that, as a result, the weak rule has a greater number of documents in the case of CFC substitutes. For HAPs, the graphs show no difference mostly due to the scale of the axis.

Table B7: Robustness check: SCM with alternative assignment rule.

(a) Patents.

Topic Means	Donor Pool	Pre RMSPE	p-value	ATE	Year
unweighted	small pool	0.37	0	1.77	1990
weighted	small pool	0.40	0	2.06	1990
unweighted	whole sample	0.40	–	1.73	1990
weighted	whole sample	0.46	0.03	1.64	1990

(b) Articles.

Topic Means	Donor Pool	Pre RMSPE	p-value	ATE	Year
weighted	small pool	0.27	0.0	1.34	1990
unweighted	small pool	0.28	0.0	1.35	1990
weighted	whole sample	0.28	0.0	1.39	1990
unweighted	whole sample	0.28	0.0	1.31	1990

Note: In this robustness check, I use an alternative rule to assign documents to molecule. The estimated treatment effects for these robustness checks are all higher than the effect estimated with the main methodology. "Topic Means" indicates the procedure for aggregating the topic proportions at the molecule level. If "weighted", the calculated proportion of topic j for molecule i is the mean proportion of topic j across all documents mentioning molecule i, weighted by the number of times the molecule appears in the document. "Donor Pool" indicates what sample of HAPs is used in the SCM procedure. For "small pool", the sample of HAPs used corresponds to the twenty HAPs most similar to the treated unit in terms of counts and topic proportions before 1987.

Table B8: Robustness check: SCM with ten topics.

(a) Patents.

Rule	Topic Means	Donor Pool	Pre RMSPE	p-value	ATE	Year
weak	weighted	small pool	0.17	0	1.16	1990
weak	unweighted	small pool	0.17	0	1.17	1990
weak	unweighted	whole sample	0.17	0.035	0.62	1991
weak	weighted	whole sample	0.19	–	0.71	1990
intermediate	unweighted	small pool	0.52	0	2.0	1990
intermediate	unweighted	whole sample	0.52	0.041	1.86	1989
intermediate	weighted	whole sample	0.59	0.053	1.77	1989
intermediate	weighted	small pool	0.67	0	2.01	1990

(b) Articles.

Rule	Topic Means	Donor Pool	Pre RMSPE	p-value	ATE	Year
weak	unweighted	whole sample	0.25	0.000	1.01	1990
intermediate	weighted	small pool	0.26	0.000	1.16	1990
weak	weighted	whole sample	0.27	0.018	1.05	1990
intermediate	unweighted	small pool	0.29	0.000	1.25	1990
weak	weighted	small pool	0.30	0.000	1.45	1990
weak	unweighted	small pool	0.32	0.000	1.02	1992
intermediate	weighted	whole sample	0.42	0.000	1.67	1989
intermediate	unweighted	whole sample	0.42	0.000	1.68	1989

Note: In this robustness check, I increase the number of topics generated by the LDA topic model from five to ten. I, therefore, use ten different topic proportions as covariates in the SCM procedure. "Topic Means" indicates the procedure for aggregating the topic proportions at the molecule level. If "weighted", the calculated proportion of topic j for molecule i is the mean proportion of topic j across all documents mentioning molecule i, weighted by the number of times the molecule appears in the document. "Donor Pool" indicates what sample of HAPs is used in the SCM procedure. For "small pool", the sample of HAPs used corresponds to the twenty HAPs most similar to the treated unit in terms of counts and topic proportions before 1987.

Table B9: HAPs contributing to the synthetic controls and their respective weights

Corpus	Patents				Articles			
Topics	5 topics		10 topics		5 topics		10 topics	
Weights	Weighted	Unweighted	Weighted	Unweighted	Weighted	Unweighted	Weighted	Unweighted
1,4-Dichlorobenzene					0.13		0.05	0.21
2,4-Toluene diamine		0.03						
3,3-Dimethoxybenzidine	0.61	0.59	0.38	0.57				
Calcium cyanamide		0.07						
Caprolactam					0.23			
Captan					0.09			
Carbonyl sulfide							0.08	
Chloroprene					0.55			
Diethyl sulfate			0.23				0.37	
Dimethyl phthalate						0.11	0.41	
Ethyl acrylate						0.56		
Ethyl chloride	0.11			0.16				
Ethylidene dichloride				0.1				
Pentachlorophenol	0.07							
beta-Propiolactone						0.32		
m-Cresol	0.06							
o-Cresol			0.33	0				0.66
o-Toluidine							0.09	0.13
o-Xylenes	0.15	0.32	0.05	0.16				
p-Cresol			0.01					

Note: The tables describe the HAPs entering the synthetic control across four SCM specifications (unweigthed or weighted means, and 5 or 10 topics). The SCM specifications always used the small pool of HAPs and log counts.

Table B10: Description of the HAPs contributing to the synthetic controls

HAPs	Description
3,3-Dimethoxybenzidine	Used as an intermediate in the production of dyes and pigments.
o-Xylenes	Used in the production of ethylbenzene, as solvents in products such as paints and coatings, and are blended into gasoline. Released into the atmosphere as fugitive emissions from industrial sources, from auto exhaust, and through volatilization from their use as solvents
o-Cresol	Used as disinfectant, preservative, and wood preservative. Mainly used as a precursor to other compounds such as herbicides and pharmaceutical intermediates. Used commercially as a disinfectant.
Ethyl chloride	Used production of ethyl cellulose, use as a solvent, refrigerant, and topical anesthetic, in the manufacture of dyes, chemicals, and pharmaceuticals. Was used in the production of tetraethyl lead, an anti-knock additive to leaded gasoline.Âăa shift to the use of unleaded gasoline has caused a drastic reduction in the amount of ethyl chloride required for the production of tetraethyl lead.
Diethyl sulfate	Used as an ethylating agent and as a chemical intermediate. Used as an accelerator in the sulfation of ethylene and in some sulfonations. chemical intermediate for ethyl derivatives of phenols, amines, and thiols, and as an alkylating agent.
Calcium cyanamide	Used as a fertilizer, defoliant, herbicide, fungicide, and pesticide; in the manufacture and refining of iron; and in the manufacture of calcium cyanide, melamine, and dicyandiamide.
Ethylidene dichloride	Primarily used as an intermediate in the manufacture of other chemicals such as vinyl chloride and 1,1,1-trichloroethane, and to manufacture high vacuum rubber. Limited use as a solvent for plastics, oils, and fats.
Pentachlorophenol	Was once one of the most widely used biocides in the United States, but it is now a restricted use pesticide. Was used as a wood preservative; Still used for the formulation of fungicidal and insecticidal solutions and for incorporation into other pesticide products.
m-Cresol	Used as disinfectant, preservative, and wood preservative. Used to produce certain herbicides, as a precursor to the pyrethroid insecticides, to produce antioxidants, and to manufacture the explosive, 2,4,6-nitro-m-cresol.
2,4-Toluene diamine	Used primarily in the production of toluene diisocyanate, which is used in the production of polyurethane.Âă It is used as an intermediate in the synthesis of dyes and heterocyclic compounds. Also used to prepare direct oxidation black, a dye for hair and furs, and to prepare dyes for leather. Other uses: enhancement of thermal stability in polyamides, fatigue resistance and dyeability in fibers, and the preparation of impact-resistant resins, polyimides with superior wire-coating properties, benzimidazolethiols (antioxidants), hydraulic fluids, urethane foams, fungicide stabilizers, and sensitizers for explosives.
Chloroprene	Used primarily in the manufacture of polychloroprene (Neoprene TM, duprene) which is a polychloroprene elastomer that is used to make diverse products including adhesives, automotive and industrial parts (e.g., belts and hoses), wire and cable covers, adhesives, caulks, flame-resistant cushioning and other applications requiring chemical, oil, and/or weather resistance.
Ethyl acrylate	Used in the manufacture of water-based latex paints and adhesives, textile and paper coatings, leather finish resins, and in the production of acrylic fibers
Dimethyl phthalate	Used in solid rocket propellants, lacquers, plastics, safety glasses, rubber coating agents, molding powders, insect repellants, and pesticides
1,4-Dichlorobenzene	Used mainly as a fumigant for the control of moths, molds, and mildews, and as a space deodorant for toilets and refuse containers. Also used as an intermediate in the production of other chemicals, in the control of tree-boring insects, and in the control of mold in tobacco seeds.
beta-Propiolactone	Used for vaccines, tissue grafts, surgical instruments, and enzymes, as a sterilant of blood plasma, water, milk, and nutrient broth, and as a vapor-phase disinfectant in enclosed spaces.ÂăIts sporicidal action is used against vegetative bacteria, pathologic fungi, and viruses. Also used as a chemical intermediate.
Caprolactam	Primarily used in the manufacture of synthetic fibers (especially Nylon 6). Also used in brush bristles, textile stiffeners, film coatings, synthetic leather, plastics, plasticizers, paint vehicles, cross-linking for polyurethanes, and in the synthesis of lysine.
o-Toluidine	Primarily used in the manufacture of dyes. It is also used in the manufacture of rubber vulcanization accelerators, hypnotic and anesthetic pharmaceuticals, and pesticides.
Captan	predominantly used in agriculture as a fungicide on a wide variety of fruits, vegetables, and ornamentals on plant seeds, and also on food crop packaging boxes. also used in cosmetics and pharmaceuticals, oil-based paints, lacquers, wallpaper paste, plasticizers, polyethylene, vinyl, rubber stabilizers, and textiles.
Carbonyl sulfide	Used as an intermediate in the synthesis of organic sulfur compounds and alkyl carbonates
p-Cresol	Used as disinfectant, preservative, and wood preservative. Used largely in the formulation of antioxidants and in the fragrance and dye industries.

Note: The tables provide a description of the various HAPs entering the synthetic control across four SCM specifications (unweigthed or weighted means, and 5 or 10 topics). The information displayed in the "Keywords" column was collected from the EPA. The SCM specifications always used the small pool of HAPs and log counts.

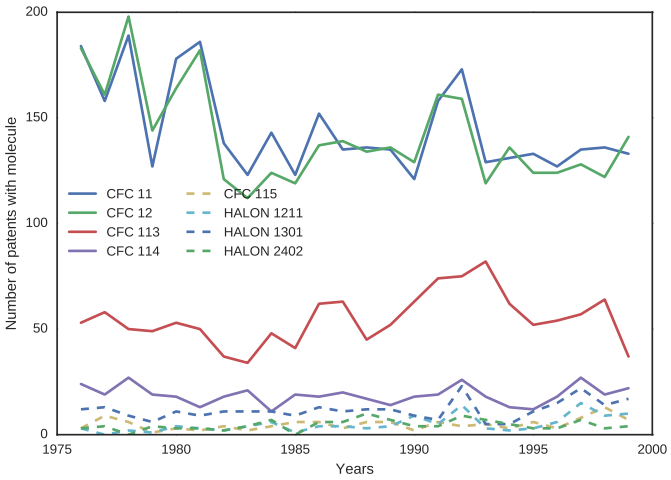

(a) Annex A: Patents

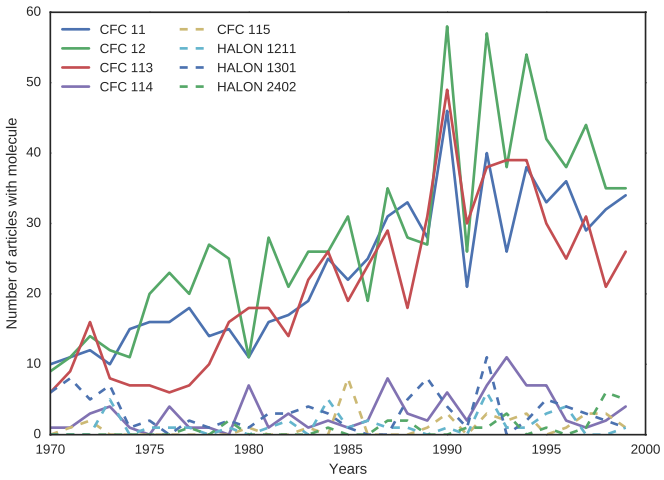

(b) Annex A: Articles

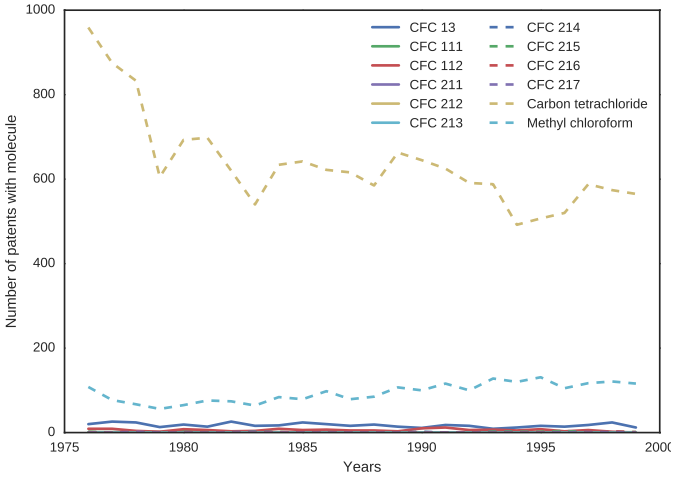

(c) Annex B: Patents

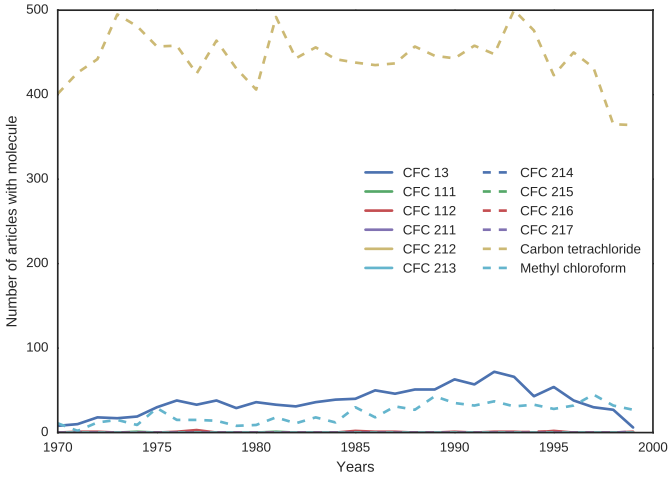

(d) Annex B: Articles

Figure B7: Counts in articles and patents for each molecule of Annex A and Annex B. Note: These graphs plot the yearly number of articles or patents mentioning the names of given molecules included in Annex A and B. We note that most trends are flat, except maybe for Annex A in articles which seem to increase and then decrease.

Table B11: First differences for Annex A and B compounds

(a) Annex A - Patents

	(1)	(2)
Post 1987	0.178***	
	(0.055)	
Post 1987 x Years		-0.009
		(0.017)
Years		0.019*
		(0.011)
Molecule FE	Yes	Yes
R-squared	0.932	0.933
Observations	184	184

Standard errors in parentheses
Dependent variable: Log count of patents
Years are relative to 1987.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

(b) Annex A - Articles

	(1)	(2)
Post 1987	0.526***	
	(0.070)	
Post 1987 x Years		-0.008
		(0.017)
Years		0.034***
		(0.008)
Molecule FE	Yes	Yes
R-squared	0.838	0.843
Observations	240	240

Standard errors in parentheses
Dependent variable: Log count of articles
Years are relative to 1987.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

(c) Annex B - Patents

	(1)	(2)
Post 1990	0.175***	
	(0.058)	
Post 1990 x Years		0.020
		(0.018)
Years		0.001
		(0.007)
Molecule FE	Yes	Yes
R-squared		
Observations	.971	.97
N	207	207

Standard errors in parentheses
Dependent variable: Log count of patents
Years are relative to 1990.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

(d) Annex B - Articles

	(1)	(2)
Post 1990	0.152***	
	(0.056)	
Post 1990 x Years		-0.040**
		(0.017)
Years		0.022***
		(0.005)
Molecule FE	Yes	Yes
R-squared		
Observations	.975	.977
N	210	210

Standard errors in parentheses
Dependent variable: Log count of articles
Years are relative to 1990.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The regressions indicate statistically significant mean shift between before and after 1987, except for Annex B in patents; however these are small in magnitude.

Table B12: Difference-in-differences for Annex A and B compounds.

(a) Annex A - Patents

	(1)	(2)
Post 1987 x Annex A	0.024	
	(0.055)	
Post 1987 x Annex A x Years		-0.009
		(0.017)
Annex A x Years		0.005
		(0.011)
Years		0.014***
		(0.001)
Year FE	Yes	No
Molecule FE	Yes	Yes
R-squared	0.987	0.985
Observations	736	736

Standard errors in parentheses
Dependent variable: Log count of patents
Years are relative to 1987.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

(b) Annex A - Articles

	(1)	(2)
Post 1987 x Annex A	0.069	
	(0.073)	
Post 1987 x Annex A x Years		-0.008
		(0.017)
Annex A x Years		0.006
		(0.008)
Years		0.028***
		(0.002)
Year FE	Yes	No
Molecule FE	Yes	Yes
R-squared	0.966	0.964
Observations	960	960

Standard errors in parentheses
Dependent variable: Log count of articles
Years are relative to 1987.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

(c) Annex B - Patents

	(1)	(2)
Post 1990 x Annex B	0.080	
	(0.063)	
Post 1990 x Annex B x Years		0.020
		(0.018)
Annex B x Years		-0.004
		(0.007)
Years		0.005***
		(0.002)
Year FE	Yes	No
Molecule FE	Yes	Yes
R-squared		
Observations	.988	.987
N	828	828

Standard errors in parentheses
Dependent variable: Log count of patents
Years are relative to 1990.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

(d) Annex B - Articles

	(1)	(2)
Post 1990 x Annex B	-0.251***	
	(0.065)	
Post 1990 x Annex B x Years		-0.040**
		(0.017)
Annex B x Years		-0.004
		(0.005)
Years		0.026***
		(0.002)
Year FE	Yes	No
Molecule FE	Yes	Yes
R-squared		
Observations	.968	.967
N	840	840

Standard errors in parentheses
Dependent variable: Log count of articles
Years are relative to 1990.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The difference-in-differences specifications indicate that a positive and statistically significant treatment effect for Annex in patents and a negative one for Annex B in articles. The magnitudes however are small.

Table B13: SCM for Annex A and B compounds

(a) Annex A - Patents

Topic Means	Donor Pool	Pre RMSPE	p-value	ATE	Year
unweighted	whole sample	0.08	0.06	-0.01	–
weighted	whole sample	0.09	0.14	-0.0	–
unweighted	small pool	0.13	0.90	-0.09	–
weighted	small pool	0.14	0.60	-0.03	–

(b) Annex A - Articles

Topic Means	Donor Pool	Pre RMSPE	p-value	ATE	Year
unweighted	small pool	0.15	0.250	-0.12	–
unweighted	whole sample	0.15	0.471	0.08	–
weighted	whole sample	0.16	0.296	0.18	–
weighted	small pool	0.25	0.400	0.33	1990

(c) Annex B - Patents

Topic Means	Donor Pool	Pre RMSPE	p-value	ATE	Year
weighted	whole sample	0.07	0.09	-0.18	–
unweighted	whole sample	0.07	0.10	-0.16	–
unweighted	small pool	0.09	0.20	-0.23	–
weighted	small pool	0.12	0.30	-0.28	–

(d) Annex B - Articles

Topic Means	Donor Pool	Pre RMSPE	p-value	ATE	Year
unweighted	whole sample	0.15	0.036	-0.32	–
weighted	whole sample	0.16	0.112	-0.28	–
weighted	small pool	0.38	0.200	-0.86	–
unweighted	small pool	0.40	0.300	-0.86	–

Note: All tables refer to SCM implementation using log count as outcome variable and using the weak rule of assigning documents. Almost none of the procedures yield treatment effects that are statistically significant. Most p-values are greater than 0.10. These results indicate that Montreal did not trigger a large decrease nor a large increase in the number of patetns and articles mentioning Annex A and B compounds. "Topic Means" indicates the procedure for aggregating the topic proportions at the molecule level. If "weighted", the calculated proportion of topic j for molecule i is the mean proportion of topic j across all documents mentioning molecule i, weighted by the number of times the molecule appears in the document. "Donor Pool" indicates what sample of HAPs is used in the SCM procedure. For "small pool", the sample of HAPs used corresponds to the twenty HAPs most similar to the treated unit in terms of counts and topic proportions before 1987.

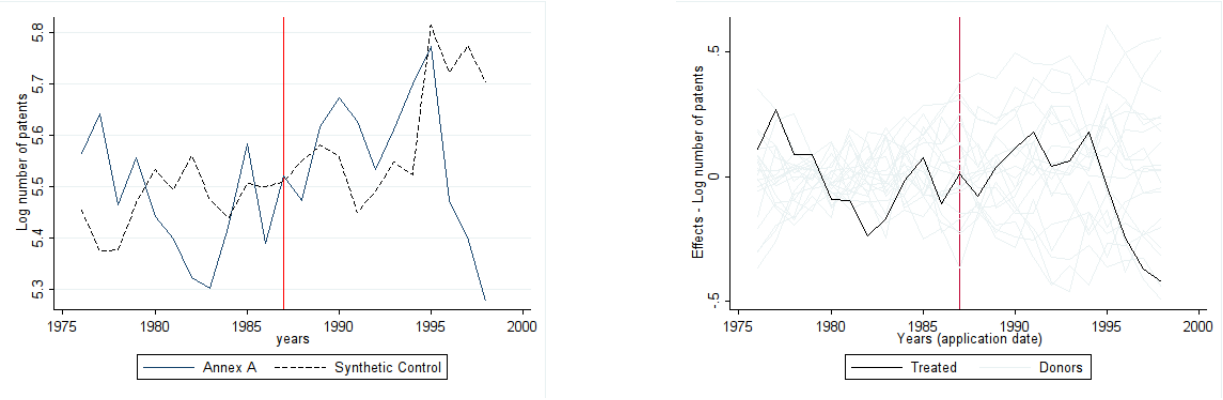

(a) Patents: raw effect (left) and placebo tests (right)

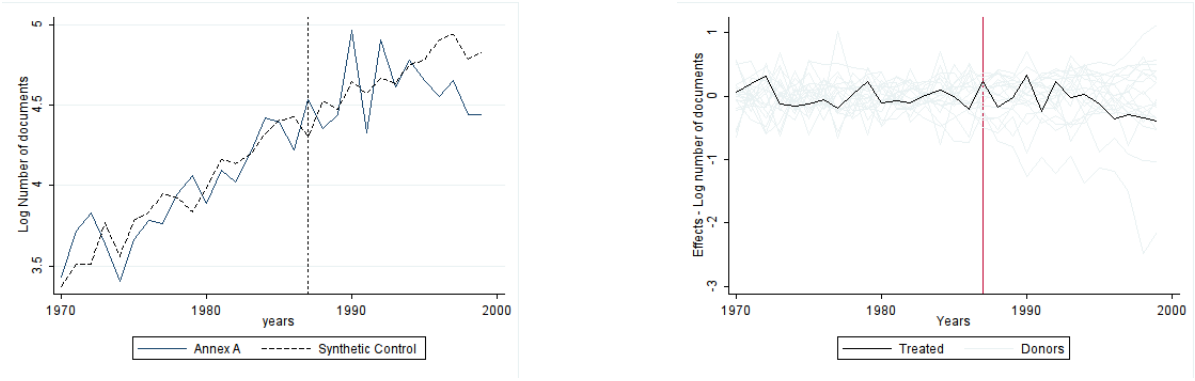

(b) Articles: raw effect (left) and placebo tests (right)

Figure B8: SCM for Annex A compounds

Note: Figures B8a and B8b display the results of the synthetic control method for Annex A compounds for patents and articles. There is no significant increase or decrease in the number of patents and articles mentioning Annex A compounds. In all cases, the method is implemented using the topic proportions of a LDA model with 5 topics and the weak rule for assigning documents to molecule groups. Weighted means of topic proportions are used for patent and unweighted means for articles because these are the specifications that yielded lowest pre-RMPSE. The graphs on the left-hand side represent the raw effect, that is the observed time series of the treated group along with the time series of the constructed control. On the right-hand sides are shown the placebo tests, the non-parametric tests to evaluate the significance of the results; black lines show the effect on the treated group relative to the control group, while each gray line is a placebo test performed on an unit drawn from the donor pool.

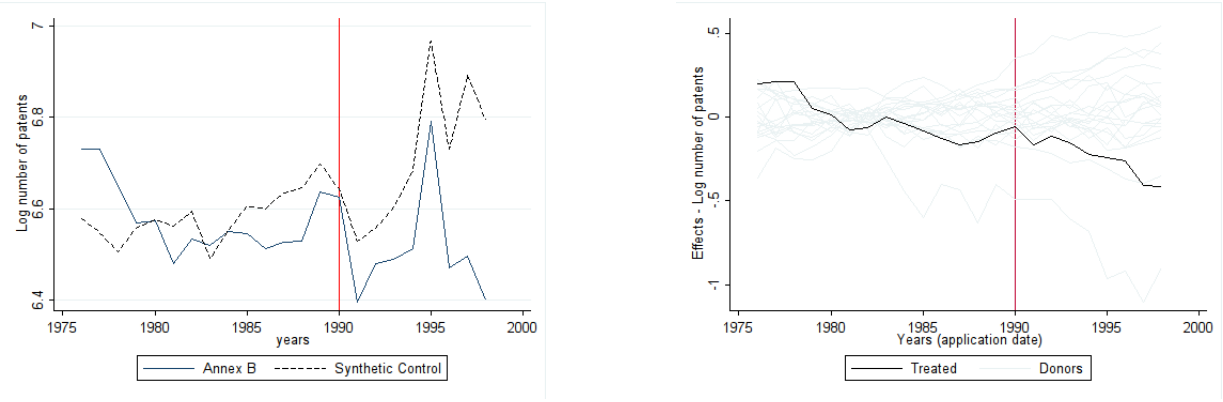

(a) Patents: raw effect (left) and placebo tests (right)

(b) Articles: raw effect (left) and placebo tests (right)

Figure B9: SCM for Annex B compounds

Note: Figures B9a and B9b display the results of the synthetic control method for Annex B compounds for articles and patents. We note that in this case the synthetic control offers a poor fit to the observed data. Hence we cannot infer whether there is anincrease or a decrease. In all cases, the method is implemented using the topic proportions of a LDA model with 5 topics and the weak rule for assigning documents to molecule groups. Weighted means of topic proportions are used for patent and unweighted means for articles because these are the specifications that yielded lowest pre-RMPSE. The graphs on the left-hand side represent the raw effect, that is the observed time series of the treated group along with the time series of the constructed control. On the right-hand sides are shown the placebo tests, the non-parametric tests to evaluate the significance of the results; black lines show the effect on the treated group relative to the control group, while each gray line is a placebo test performed on an unit drawn from the donor pool.