

Induced Innovation and
International Environmental
Agreements: Evidence from
the Ozone Regime

Eugenie Dugoua

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Induced Innovation and International Environmental Agreements: Evidence from the Ozone Regime

Abstract

This paper revisits one of the rare success stories in global environmental cooperation: the Montreal Protocol and the phase-out of ozone-depleting substances. I show that the protocol increased science and innovation on alternatives to ozone-depleting substances, and argue that agreements can indeed be useful to solving global public goods problems. This contrasts with game-theoretical predictions that agreements occur only when costs to the players are low, and with the often-heard narrative that substitutes were readily available. I reconcile theory and empirics by discussing the role of induced innovation in models of environmental agreements.

JEL-Codes: O300, Q550.

Keywords: induced innovation, directed technological change, green innovation.

Eugenie Dugoua
London School of Economics
Houghton St, Holborn
United Kingdom – London WC2A 2AE
e.dugoua@lse.ac.uk

June 18, 2023

I received no specific funding for this work. I am deeply grateful to Scott Barrett, Geoffrey Heal, and Suresh Naidu for their advice and support. I thank Timothy Bresnahan, David Popp, Bhaven Sampat, Glenn Sheriff, Max Auffhammer, Matthew Kotchen, William McAllister, Juan Moreno-Cruz, Roger Fouquet, Cristian Pop-Eleches, Bernard Salanie, Johannes Urpelainen, Doug Almond, Wolfram Schlenker, Rodrigo Soares, Johannes Boehm, Marion Dumas, Anna Thompsett, Kyle Meng, Prabhat Barnwal, Jeffrey Shrader, Anthony D'Agostino, Eyal Frank and Anouch Missirian for their helpful comments, as well as the participants of the NBER Innovation Summer Institute and the many seminars where this paper was presented and discussed. All errors are mine.

1 INTRODUCTION

Global collective action problems are some of the most pressing issues that humanity is facing. Environmental concerns such as climate change or biodiversity have seen minimal progress, but one shining example of success is the fight against ozone depletion. At the end of the 1970s, scientists warned that chlorofluorocarbons (CFCs) might destroy ozone molecules in the stratosphere and reduce the extent to which humans were protected from solar radiation. The issue rose to the top of the global agenda, and, in 1987, high-income countries negotiated the Montreal Protocol to phase out CFCs from industrial activities.

In Montreal, signatories agreed to decrease the production and consumption of CFCs following a schedule of reduction targets. The protocol included trade restrictions with non-parties in ozone-depleting substances and threats of banning trade in products made using ozone-depleting substances. Technological change unfolded rapidly, and within a decade, the production and consumption of CFCs decreased by more than 80%.¹ Such success story presents us with a formidable opportunity to study what worked.

Although the Montreal Protocol remains a point of reference in discussing global environmental problems (Barrett 1999; Sunstein 2007), the dynamics of innovation in the ozone crisis are still debated. Some, like Richard E. Benedick, the chief U.S. negotiator at Montreal, argued the agreement triggered a vast effort in research to find CFC substitutes². Others claim that CFC substitutes were already available at the time of negotiations or that the industry endorsed CFC cuts because it had achieved a breakthrough (Heal 2016; Sunstein 2007).

This paper offers the first quantitative study of whether the Montreal Protocol induced science and innovation on CFC substitutes. To do so, I compile a list of 14 molecules that scientists and industry experts identified as best candidates for CFC substitutes and consider those molecules as treated by the Montreal Protocol. I then track mentions of CFC substitutes over time in scientific articles published in journals indexed by ScienceDirect and patents granted by the United States Patent and Trade Office (USPTO).

The primary hypothesis is that the Montreal Protocol provided a clear signal and power-

1. My calculations using UNEP data.

2. "It was evident (...) that the protocol was in fact moving industry in directions that two years earlier had been considered impossible." (Benedick 2009, Chap.8 p.104.)

ful incentives for firms and scientists to increase work on CFC substitutes, which led to an increase in patents and scientific articles mentioning these molecules. The empirical strategy relies on comparing CFC substitutes with a control group of molecules used in industrial applications similar to CFC substitutes but with no connection to ozone, known as HAPs (Hazardous Air Pollutants). To ensure that the molecules are comparable, I use topic modeling algorithms on the text of patents and articles to construct molecule-level variables that proxy for the molecules' scientific and industrial context.

As illustrated in Figure 1, only a few patents and articles on CFC substitutes were published before 1987, and the trend before 1987 is flat, possibly indicating that the agreement was little anticipated. The difference-in-differences (DiD) suggests that the protocol led to a 400% increase in the number of patents related to CFC substitutes (relative to the pre-treatment period) and a 500% increase in the case of scientific articles. The increase becomes statistically significant, starting in 1989 for patents and 1990 for articles, two to three years after the agreement's signature.³ The estimates are robust to a series of alternative specifications, including weighting counts by the number of occurrences of the molecule's name in the text and weighting by the number of citations that the document received. As an alternative approach, I also estimate the protocol's impact using a synthetic control method and find consistent, yet smaller, increases of about 135% for patents and 180% for articles.⁴

A possible alternative story is that firms undertook research and development before the negotiations and kept substitutes "hidden" by not filing patents until the agreement was announced. I argue that if firms kept hidden some of their innovation, we should expect a one-time increase in patent counts in the immediate aftermaths of Montreal. Indeed, patenting was the primary mechanism firms had to protect their innovations. On the one hand, trade secrets were very hard to keep in the development of CFC substitutes (Parson 2003). On the other

3. This delay is similar to prior results in the literature (Popp 2002) and can be attributed to the time required to turn research efforts into patent applications and published academic papers.

4. Shapiro and Warhit (1983) estimated that limiting CFC production to 1980 levels would lead to price increases between 500% and 1000% depending on the CFC compounds. Given this, a 135% increase in patents corresponds to an elasticity between 0.14 and 0.27. These estimates are broadly consistent with related studies in different contexts, such as Popp (2002), who reports short-run and long-run energy price elasticities of 0.060 and 0.354, respectively (see Table 2 in the article) and Aghion et al. (2016), who estimate the elasticity of clean car patents with respect to gasoline prices to be between 0.843 and 0.970, depending on the specification (see Table 3 in the article).




FIGURE 1

Counts of Patents and Articles Mentioning CFC Substitutes

Note: The graph plots the yearly number of patents and articles mentioning the names of any of the 14 CFC substitutes. The x-axis corresponds to the *application* date of patents granted between 1976 and 1999 and the *publication* date of articles. We note that only a few patents and articles on CFC substitutes were published before 1987 and that the trend up to 1987 is flat. After 1987, we observe a large increase in both patents and articles mentioning CFC substitutes. A version of this figure with trends for CFC compounds is shown in Online Appendix Figure I3.

hand, the relatively high degree of competition in the industry ensured that several firms were working on related technologies, and delaying patenting meant a higher risk of competitors patenting first. Assuming firms did not patent their innovation before 1987, the signature of Montreal would have acted as a strong positive shock to their incentives to seek intellectual property rights. I find that patenting activity shows no spike after the signing of the agreement but, instead, a progressive ramp up. This lends little support to the narrative that the industry had achieved some breakthrough which they would have kept secret.

The nature of the technological challenge was, in fact, not consistent with the idea of a breakthrough. First, some CFC reductions were easy to achieve, and in particular, significant cuts could be achieved simply by using CFCs more efficiently or with recovery and recycling. Some chemical substitutes also looked promising for specific applications. For example, in 1988, the FDA approved HCFC-22 for fast food foams and grocery display packaging. But, in other sectors, such as the growing refrigeration and air-conditioning markets, a full phase-out required new molecules. Consequently, from a 1987 perspective, the technological problem of reducing CFC emissions could be likened to a marginal abatement cost curve that initially had a low and flat slope for a significant amount of reductions but then gradually became steeper and with significant uncertainty.⁵

Furthermore, developing CFC substitutes was not about “new-to-the-world” compounds but rather about “new-to-the-industry” compounds that required research efforts to adapt them to many industrial applications. Indeed, chemical engineers knew that the compounds with the best potential had to present molecular structures similar to CFCs. This would ensure that their thermodynamic properties would best mimic those of CFCs and make them easier to substitute in the myriad of industrial applications that used CFCs (Pool 1988). Although these molecules had been known on paper for decades, scientists still needed to learn more about their thermodynamic properties, toxicity profile, and environmental acceptability. Firms had to experiment with new processes and formulas to retrofit installed equipment with CFC substitutes or replace them altogether.

5. The availability of cheap reductions via recycling, efficiency measures, and specific chemical substitute applications explains why CFCs production decreased rapidly after 1987. For more information, see Online Appendix Figure A2 and Table A4.

I also investigate the possibility that mounting consumer pressure played a role in fostering science and innovation on CFC substitutes.⁶ Already in 1985, scientists showed that extensive depletion of ozone was taking place over Antarctica. The so-called ozone “hole” was causally attributed to CFCs in March 1988, likely increasing the saliency of the CFC threat in the public’s eyes. To further investigate the role of consumer pressure on fostering innovation on CFC substitutes, I collect data from the EPA website on all industrial applications of CFC substitutes and classify substitutes as either consumer exposed, not consumer exposed, both or undetermined. I find that the increase in patenting for molecules that were not used in consumer appliances is similar to the increase seen in the broader sample of molecules, suggesting that consumer pressure is unlikely to have driven the large increase in patenting.

At first sight, it may seem counterintuitive that the Montreal Protocol induced innovation, or more generally, did anything that would not have happened without it. As shown in Barrett (1994), theoretical models make the dismal prediction that international environmental agreements occur only when cooperation is easy to achieve, that is, when costs to the players are low and benefits high. This implies that the Montreal Protocol could not have been far from what countries would have done unilaterally. I revisit this interpretation by discussing how induced innovation can affect models of international environmental agreements. With induced innovation, marginal abatement costs decrease over time as emission reductions are undertaken. This allows for more ambitious targets to be agreed upon, even though they would have been outside of negotiation reach early on in the process.

Hence, the signature of the Montreal Protocol should be seen as a first modest step in a series of increasingly ambitious agreements. The targets agreed in 1987 required a 50% cut by 1998, while those negotiated in the London and Copenhagen amendments (in 1990 and 1992, respectively) enforced a complete ban and added a broader range of molecules to the list of regulated ozone-depleting substances. The targets agreed in 1987 may be interpreted as modest because they encoded what the industry thought feasible over a reasonable period. This did not mean the industry already had all the necessary knowledge and technologies. But it meant that they expected that they could meet the targets with sufficient investments in research and

6. Yet another possibility could be the existence of research grants or subsidy programs specifically financing research and development work on CFC substitutes. To my knowledge, no such programs were ever implemented.

development. Notably, the agreement ensured a leveled playing field for all firms competing in the CFC industry. Indeed, contrary to many other agreements, a key aspect of the Montreal Protocol is that it built in an enforcement mechanism. Montreal included trade restrictions, and threats of banning trade in any products made using CFCs, making it economically binding.

Agreements close to the non-cooperative equilibrium, like Montreal, should not be seen as agreements where costs to the players are low, but rather as agreements where costs are *expected* to be low. Between expectation to realization come the actual efforts of doing research and development, of implementing ideas that exist only on paper and that still require experimentation and tinkering. Such endeavors improve technologies making them more affordable and increasing the scope of their uses. By inducing firms to innovate, Montreal reduced the expected cost of further emission abatement. In turn, this made it easier for governments and industries to negotiate binding amendments with more ambitious targets.

This paper shows that the ozone layer's success story, therefore, is better summarized as a series of agreements that progressively ramped up ambitions in emission reductions. We may think of it as a repeated cooperation game where, at each stage, small but binding reductions force firms to innovate. This reduces expected abatement costs and leads to better cooperation outcomes in the next stage. Innovation here plays a critical role in enabling ambitions to ratchet up. By showing that a low-ambition but binding agreement such as the Montreal Protocol did encourage the development of technological solutions, this paper suggests such agreements are potent tools that dynamically improve the benefit-cost balance of environmental protection and may therefore also be useful to deal with current problems such as climate change.

This paper contributes to the literature on technological change and the environment (Jaffe, Newell, and Stavins 2002; Popp 2019; Popp, Newell, and Jaffe 2010). Recent studies have drawn attention to the factors inducing innovation in environmental-friendly technologies. In particular, prior research has emphasized both the role of energy price increases (Aghion et al. 2016; Popp 2002) and domestic environmental policies (Calel and Dechezleprêtre 2016; Jaffe and Palmer 1997). This paper, instead, documents that agreements for global environmental public goods can induce innovation. In principle, like domestic policies, agreements should force firms to allocate inputs to reduce pollution, thereby inducing them to innovate

(Milliman and Prince 1989; Porter and Van Der Linde 1995). Several studies investigate the effect of such agreements on pollution outcomes (Aichele and Felbermayr 2011; Finus and Tjøtta 2003; Kellenberg and Levinson 2014), but they seldom look at the impact on science and innovation.⁷

This paper also contributes to the literature on the economics of international environmental agreements (Barrett 1994; Battaglini and Harstad 2016; Harstad, Lancia, and Russo 2019) and ozone cooperation (Auffhammer, Morzuch, and Stranlund 2005; Barrett 1994, 2003; Murdoch and Sandler 2009). In particular, Barrett (1999) suggested Montreal’s trade measures solved the enforcement problem, and Wagner (2016) further argued they promoted full participation in the protocol, ensuring its almost-universal ratification. Outside of economics, research has examined aspects relating to negotiations and diplomacy (Andersen and Sarma 2012; Benedick 2009) as well as qualitative accounts of corporate strategy and innovation (Falkner 2005; Mulder 2005; Parson 2003; Smith 1998; Taddonio, Sarma, and Andersen 2012).

The paper first provides background information in Section 2. Section 3 then describes the data, Section 4 outlines the empirical strategy, Section 5 presents the main results, and Section 6 considers mechanisms and alternative hypotheses. Finally, Section 7 discusses the role of induced innovation in theoretical models of global environmental cooperation to reconcile the empirical results with the literature on international environmental agreements. Section 8 concludes.

2 BACKGROUND

2.1 Theory and Hypotheses

When factors of production become more expensive, technological change tends to be biased towards them, so to make their use more efficient or substitute them. Such theory of *induced innovation* was initially developed by Hicks (1932), but the concept has reappeared in the past two decades under the phrase “directed technical change,” encompassing not just price effects but also market size and regulatory effects (Acemoglu 1998). The idea was later augmented to

7. One exception is Dekker et al. (2012) who focuses on transboundary air pollution.

include environmental policies. In the simplest model, environmental regulations force firms to allocate inputs (labor or capital) to pollution reduction or restrict the choice of technologies and inputs in the production process (Milliman and Prince 1989). As such, environmental regulations are modeled as extra costs.

The theory of *induced innovation* argues that environmental policies induce firms to innovate in the hope of offsetting regulation-imposed costs, at least partially.⁸ The reasoning can easily be extended to international environmental agreements. Agreements oblige firms to adapt their production processes. The signature on its own immediately changes expectations regarding future domestic environmental policies, as long as firms perceive the agreement as binding. Arguably, this is what happened with the Montreal Protocol.

Montreal included trade restrictions with non-parties in products containing those ozone-depleting substances and a threat of banning trade in products made using ozone-depleting substances. The trade restrictions effectively acted as a mechanism for free-rider deterrence and leakage prevention, rendering the agreement binding. Therefore, the main hypothesis is that Montreal credibly signaled firms that the continued use of CFCs would become increasingly costly, and induced them to innovate. Richard Benedick, the U.S. head negotiator at Montreal, argued that “(it) was evident (...) that the protocol was in fact moving industry in directions that two years earlier had been considered impossible” (Benedick 2009, Chap.8 p.104.). However, Benedick only refers to articles published in the *New York Times* and *Chemical and Engineering News* to support his claim.

On the other hand, an often heard narrative argues that CFC substitutes were readily available before the negotiations. Sunstein (2007) claims that “an international agreement was largely in the interest of American manufacturers, which had already initiated a transition to safe CFC-alternatives.” This view is also often expressed in media outlets. For example, *The New York Times*, on August 20, 2002, stated that “(the) agreement’s success occurred, in large part, because substitutes for the harmful chemicals were readily available (...).” Importantly, such claims are consistent with the theory of why and when agreements appear (Barrett 1994). In the case of Montreal, that theory implies that the agreement was successfully negotiated

8. The *Porter Hypothesis* goes further arguing the extra costs imposed by environmental policies can sometimes be even more than fully offset (Ambec et al. 2013; Porter and Van Der Linde 1995).

because the costs of doing so were low relative to benefits.

This paper relies on quantitative analysis to examine these different hypotheses. Section 7 will further discuss how to reconcile the empirical results with the theory of international environmental agreements.

2.2 Events That Led to Montreal

The story of the ozone crisis began in 1974 when two chemists laid out the theoretical possibility that CFCs broke down ozone molecules in the stratosphere (Molina and Rowland 1974). The harmful effects of a thinner ozone layer were not well understood, but it was clear that more UV light would cause more skin cancers, eye cataracts, and, likely, lower productivity in fishery and agriculture (Miller and Mintzer 1986).

In the late 1970s, the issue began to take prominence in the media and policy circles. A few countries and firms unilaterally decided to take action. In August 1977, the U.S. Congress wrote into law a CFC ban on aerosols by 1978,⁹ and firms such as DuPont removed CFCs from their spray products because they worried about their public image. These pre-Montreal domestic regulations and corporate decisions targeted a particular industrial application of CFCs for which cheap physical substitutes existed (e.g., pump-action sprays instead of aerosol sprays). In essence, these were zero-cost unilateral moves that did not require significant research efforts. The low numbers of patents and articles between 1970 and 1987 in Figure 1 indicate that neither aerosol regulations nor consumer pressure seemed to have stimulated science and innovation on the 14 CFC substitutes I consider in that period.¹⁰ In 1980, the EPA proposed to freeze other uses beyond aerosols, but U.S. industry blocked the initiative.¹¹

At the beginning of the 1980s, concerns over ozone depletion waned. Uncertainties in the science of atmospheric ozone seemed irreducible, and the year 1981 saw the inauguration of a strongly anti-regulatory American administration. In Europe as well, many governments persisted in opposing environmental regulations that would harm manufacturers. Parson (2003)

9. Similarly, in 1978, Canada, Switzerland and Scandinavian countries all banned CFC aerosols. Germany called for a European Community-wide ban, but without success.

10. These 14 CFC substitutes were targeting foams, refrigeration and solvent applications of CFC.

11. Regulations for ozone-depleting substances were approved by the US Congress only in 1990 via the Clean Air Act amendment.

provides a detailed qualitative account of firms' reactions during this period. Although some manufacturers initially started research on potential substitutes in the late 1970s, these efforts quickly came to an end around 1981. Manufacturers stopped such R&D programs because they estimated that CFC substitutes would cost around two to five times more than CFCs. It made no sense to continue working on these substitutes with little sign of regulations underway.¹²

The political context of Montreal's negotiations was not without its surprises. In his account of the diplomatic efforts, Benedick (2009) emphasizes the great uncertainty of the negotiations' outcome until the last minute and argues that some exceptional turns of events unlocked the situation. In particular, Reagan unexpectedly overruled his administration and approved the agreement. The U.S. President had skin cancer removed twice in the past, and it has been suggested that Reagan's life experiences weighed heavily on his decision. On the European side, the most prominent opponent to CFC regulations, the United Kingdom, left the European Community Presidency, leaving Germany, Denmark, and Belgium, firm proponents, as the head negotiators. Importantly, in this account of the negotiations, the agreement succeeded independently from the state of R&D activities on CFC substitutes.

2.3 Which Molecules Were “Treated”?

CFCs are a group of molecules with a particular structure: they contain only carbon, chlorine, and fluorine atoms. This structure drives their particularly attractive thermodynamic properties: they are unusually stable, nonflammable, nontoxic, and noncorrosive. Initially, CFCs somewhat embodied the miracle of modern chemistry as they were ideal for manufacturing many consumer goods. 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). Best of all, they were cheap to produce, and so they became broadly used in many different industries such as foams, refrigeration, air-conditioning, aerosols, fire protection, and solvents. CFCs are great refrigerants because they vaporize at low temperatures and are very energy-efficient coolants. As aerosols, they were used in cosmetics, household products, pharmaceuticals, and cleaners. Their nonreactive property also made them

12. For details, see Parson (2003, Chap.3 p.53 and Chap.7 p.173).

key products for cleaning microchips and telecommunication equipment.

Strategies for reducing CFCs included physical substitutes (like pump-action sprays instead of aerosol sprays) or recycling.¹³ However, the most critical applications, such as air-conditioning units, needed chemical substitutes. The intricate relationship between molecular structure and industrial properties implied that the set of possible substitutes was limited: good candidates required a molecular structure similar to CFCs but with fewer chlorine atoms. Such compounds are known as hydro-chlorofluorocarbons (HCFCs) and hydro-fluorocarbons (HFCs). For example, CFC-12 is a compound that contains two chlorine atoms. When a hydrogen atom replaces one chlorine, we obtain a potential CFC substitute called HCFC-22. When hydrogens replace the two chlorines, we obtain another potential CFC substitute called HFC-32.¹⁴

It was public knowledge that the quest for CFC substitutes lay in the realm of HCFCs and HFCs. These molecules had been known for a long time, at least on paper and in the lab.¹⁵ Hence, developing CFC substitutes was not so much about "new-to-the-world" compounds but instead about "new-to-the-industry" compounds. The key technological challenges lay in making large-scale production cost-efficient, redesigning processes and equipment already installed, and learning about environmental acceptability and human toxicity.¹⁶

I compile a list of potential substitutes using historical records. In December 1988, a report was issued to investigate the atmospheric dynamics of 12 potential CFC substitutes. Importantly, this report, known as the Alternative Fluorocarbon Environmental Acceptability Study

13. Montreal did not impose a total ban on CFCs which meant that CFCs would still be used, even in high-income countries; in low- and middle-income countries, their use was not yet regulated. The announcement of the freeze and progressive phase-out probably encouraged firms to cease developing new applications for CFCs. However, it may also have prompted them to develop more efficient ways of using CFCs or methods for recycling, which could translate into new patents and articles. Therefore, the impact of the Montreal Protocol on patenting related to Annex A and B compounds is inherently uncertain. It may have led to more patents for certain applications and fewer for others, making it challenging to establish a strong prior on the net effect.

14. More details are available in Online Appendix Figure A1.

15. 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. 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.

16. Two potential substitutes were already in production: HCFC-22 at a large scale and HCFC-142b at a small scale. For this reason, they were considered as potentially the cheapest and fastest substitutes. Unfortunately, HCFC-22's toxicity severely limited its applications outside of foams, and HCFC-142b was not considered for refrigerant applications because its thermodynamic properties were too different and would have required changes in equipment and processes.

(AFEAS), did not cover aspects of the molecules related to industrial activities but merely how the molecules may interact with human health and the environment once released in the atmosphere. In this paper, I include the 12 compounds studied in the AFEAS report as well as two other possible CFC substitutes mentioned in Benedick (2009) and Parson (2003).¹⁷

3 DATA DESCRIPTION

3.1 Patents

I follow prior literature in using patent counts as a proxy for innovation.¹⁸ In the chemical industry, patenting is an essential way of protecting competitive advantage from new products and processes (Sampat 2018). Since chemicals can often be “reverse engineered,” secrecy offers a limited mode of appropriation and strong incentives exist to use patenting either to protect inventions from being copied or to prevent competitors from patenting related inventions (W. M. Cohen, Nelson, and Walsh 2000).¹⁹

I collect the texts contained in the abstract and summary description of USPTO patent grants published between 1976 and 2000.²⁰ The cleaning procedure involves a series of standard steps such as lowercasing or removing punctuation.²¹ Patents contain the names, addresses, and affiliations of inventors and assignees, which I categorize by type (e.g., business, education, or government). 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. More details about how the meta-data is cleaned, matched, and classified by type are provided in the Online Appendix A. I use the OECD Citations database File to obtain data on the number of forward citations received by each patent.²² Finally, I sort patents by application date as

17. HFC-245fa and HFC-365mfc are mentioned as possible substitutes in foams. Online Appendix Table A2 shows the name and additional information about all molecules considered in the analysis.

18. Examples and reviews include Hall and Jaffe (2012), Henderson, Jaffe, and Trajtenberg (1998), Popp (2005), and Williams (2013, 2017).

19. See for example Moser (2012) for an examination of how the publication of the periodic table in 1869 made chemicals easier to reverse engineer and led chemical inventors to shift from secrecy to patents in the mid-nineteenth century.

20. The full-text patent data is available at bulkdata.uspto.gov/ It represents a total of 2,605,925 patents.

21. Full details are provided in Online Appendix A.

22. OECD, Citations database, February 2019

opposed to the date when they are granted to use a better measure of when the ideas present in the patents initially emerged.

3.2 Articles

The development of CFC substitutes required a better understanding of the fundamental thermodynamic properties of the proposed molecules and how they would interact with humans and the environment. Following prior literature, I use counts of scientific articles as a proxy for increases in such knowledge.²³ I collect scientific articles published between 1970 and 2000 in journals indexed by ScienceDirect, which hosts articles from about 2,500 academic journals published by Elsevier. The sample provides excellent coverage of journals in fields related to engineering and physical sciences and allows to characterize trends in the emergence of scientific knowledge related to CFCs substitutes. I used ScienceDirect's API to download the full text of articles in journals from the following disciplines: chemistry, chemical engineering, engineering, environmental science, materials science, and physics and astronomy.²⁴ After a series of cleaning procedures described in Online Appendix A, I obtain a total number of 1,811,301 articles. For data on affiliations and citation counts, I query the Scopus search API, and use the Global Research Identifier Database to classify authors' affiliations (e.g., education or company).²⁵

3.3 Searching for Molecule Names

Chemical compounds often go by several names; for example, HCFC-22 has 39 other possible names, such as chlorodifluoromethane or algeon 22. I develop an automatic script to collect all possible names on SciFinder, a database of chemical information maintained by the American Chemical Society, and search through the text of patents and articles for any occurrence of

23. Examples and reviews include Azoulay, Graff Zivin, and J. Wang (2010), Pierre Azoulay, Fons-Rosen, and Zivin (2019), Iaria, Schwarz, and Waldinger (2018), Redner (2005), Thompson and Fox-Kean (2005), and D. Wang, Song, and Barabási (2013).

24. Journals are listed by disciplines on Elsevier's website: <https://www.elsevier.com/solutions/sciencedirect/content/journal-title-lists>.

25. <https://www.grid.ac/>

these names.²⁶ The frequency with which molecules are mentioned in any one document can vary widely. As a robustness check, I construct a measure of counts weighted by the number of occurrences of molecules in the documents in Section 5.2.

In total, I find 3270 patents and 1926 articles mentioning at least one CFC substitute. Online Appendix Tables G1, G2, and G3 illustrate the types of patents and articles that mention CFC substitutes. The most common patent codes are related to chemical compounds containing halogen atoms. The most cited articles focus on physical characteristics or new synthesis routes, while the most cited patents correspond to innovations in the pharmaceutical sector. Indeed, CFCs were essential components of aerosol delivery systems for certain medications. For instance, the most frequently cited patent pertains to an aerosol formulation incorporating P134a and salbutamol, a drug commonly used to alleviate asthma symptoms and chronic obstructive pulmonary disease.²⁷ More than 96% of patents are granted to for-profit organizations, while the rest is filed by educational and public sector organizations. American assignees represent about 60% of patents, European around 25%, and Japanese around 12%.²⁸

4 EMPIRICAL STRATEGY

4.1 A Sharp Post-1987 Increase

In this section, I examine the temporal trends in the yearly count of documents about CFC substitutes. As Figure 1 illustrates, the numbers of patents and articles increase after the signature of Montreal in 1987.²⁹ I quantitatively investigate these temporal patterns with first-difference specifications. Equation 1, below, models a mean shift while Equation 2 models a trend-break.

26. 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. A full list of all the possible names of CFC substitutes is shown in Online Appendix Table A6.

27. The prevalence of pharmaceutical applications in highly-cited patents may not necessarily indicate superior patent "quality" compared to inventions in other sectors. Instead, it may simply reflect the industry's tendency to produce and cite numerous patents.

28. Online Appendix Table G4 displays summary statistics about countries and affiliations of patent assignees and authors of articles.

29. Online Appendix Figure C1 plots time-series similar to Figure 1, but each CFC substitutes separately. The post-1987 increase is present for most substitutes.

$$\text{Count}_{mt} = \alpha + \beta_0 \times \lambda_{post1987} + \lambda_m + \varepsilon_{mt} \quad (1)$$

$$\text{Count}_{mt} = \alpha + \beta_1 \times \text{Years} \times \lambda_{post1987} + \beta_2 \times \text{Years} + \lambda_m + \varepsilon_{mt} \quad (2)$$

Count_{mt} is the number of documents mentioning molecule m in year t ; $\lambda_{post1987}$ 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. The main hypothesis is that β_0 and β_1 are both positive for CFC substitutes, implying significant increases in research and patenting activities relating to CFC substitutes after 1987 once Montreal passed.

Table 1 presents the results for these specifications. The sample here consists of the 14 different CFC substitutes for which I track the number of patents and articles throughout the years. I run separate regressions for patents and articles, and bootstrap standard errors. Model 1 confirms that a significant and positive mean shift after 1987 in the number of patents and articles mentioning CFC substitutes. The coefficients indicate almost 30 additional patents and around 13 additional articles for the average CFC substitute every year after 1987. This corresponds to a 551% increase (594% increase) in the number of patents (articles) for the years 1987-2000 relative to the years 1975-1986. 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 articles.

The post-1987 increase, however, may be driven by factors other than Montreal (e.g., policies or macroeconomic conditions fostering academic and industrial research in the 1990s). To further investigate whether the post-1987 increase can be causally attributed to the protocol, I use a set of molecules known as HAPs as a comparison group.

4.2 Hazardous Air Pollutants as a Comparison Group

HAPs is an umbrella term for molecules categorized as *hazardous air pollutants* that became monitored under the 1990 Clean Air Act due to adverse ecological impacts and human health concerns including cancer, asthma, congenital disabilities, reproductive effects, and neurode-

developmental effects. Examples include benzene, chromium, or formaldehyde.³⁰

HAPs are potentially an appropriate comparison group for CFC substitutes because they are used in a multitude of industrial applications. I provide more evidence of the similarity between HAPs and CFC substitutes in Subsection 4.4 using patent codes and topic proportions.³¹ Importantly, HAPs are also unrelated to ozone depletion and are therefore not affected by the Montreal Protocol.³²

One possible concern, however, is that research efforts on CFC substitutes and HAPs are substitutes to each other and that research on CFC substitutes crowds out research on HAPs. This would violate the Stable Unit Treatment Value Assumption (SUTVA), which requires that the response of the treated unit depends only on the treatment assigned to that unit and that there are no spillovers between units in different treatment groups. Comparing patents assignees of CFC substitutes and HAPs, I find that 75% of CFC substitutes assignees never patented on HAPs.³³ Few firms are active in both CFC substitutes and HAPs, which supports the SUTVA.³⁴

Another concern is that other policies may have impacted research and innovation related to HAPs during the same period of analysis. In particular, in 1990, an amendment to the Clean Air Act required the EPA to promulgate regulations establishing emission standards for large sources of HAPs. However, since the EPA only published the initial promulgation schedule in 1993, HAPs are not, in practice, impacted by this policy change until later in the period that I consider.³⁵ I nonetheless take a conservative approach and limit my analysis to the time-period

30. The full list of HAPs is displayed in Online Appendix Table A5.

31. To collect patents and articles related to HAPs, I proceed in the same way as for CFC substitutes. The search procedure may lead to measurement error with over-detection if documents mention molecules that are not core to the subject of the document, but there is no reason to believe that the measurement error rate would differ for treated and control molecules.

32. Three HAPs (chlorine, methylenechloride, trichloroethylene) were suggested as possible CFC substitutes. I therefore exclude them from the sample. I use the EPA SNAPs website to obtain the list of compounds or devices suggested by firms to the EPA as possible alternatives to CFCs.

33. For this exercise, I focused on the top HAPs entering the synthetic control.

34. Examples of assignees that file patents related to both CFC substitutes and HAPs include 3M, Allied Chemical, BASF, Dow Chemical, and Procter & Gamble. These are all very large and broad companies from the chemical industry which likely would host CFCs and HAPs activities in different business units.

35. The EPA published the initial list of "source categories" in 1992 (i.e. the list of industries and production processes targeted by the regulations), and in 1993, the promulgation schedule specified by which year sectors were expected to comply with the emission standards for each category or subcategory of major sources and area sources of HAPs. This was known as the National Emission Standards for Hazardous Air Pollutant (NESHAP), and most sectors were asked to comply by 1997 or 2000.

until 1992 for patents and 1995 for articles. The additional three-year period for articles is to account for further delays between submission and publication of scientific articles.³⁶

4.3 Topic Modeling

I use a Latent Dirichlet Allocation (LDA) algorithm to model which *topics* are present in the documents (Blei 2012; Blei and Lafferty 2009). In this context, a topic means a distribution over words, and a document a distribution over topics. The *number* of topics is a parameter chosen by the experimenter. I run several LDA models, each with a different number of topics, and compute their coherence score (Röder, Both, and Hinneburg 2015).³⁷

For each corpus, I choose the lowest number of topics that offers the highest coherence score, which is 20 for patents and 15 for articles.³⁸ Online Appendix Tables B1 and B2 provide the full list of words the topics contain. The documents are highly technical, and topics can be challenging to interpret. For example, for patents, the most likely word in Topic 1 is “polymer” (with probability 0.61). Then come “catalyst” and “carbon.” A trained chemist may suggest labeling Topic 1 “Catalysts on polymer substrates.”³⁹

I then use the trained topic models to generate document-level variables called topic proportions. These variables take values between 0 and 1 and indicate to what extent a topic is present in a document. Said differently, the model uses the words appearing in a document to infer the proportion of each topic in that document. I aggregate topic proportions at the molecule level by calculating weighted means where the weights are proportional to the number of times a document mentions a molecule.⁴⁰ The molecule-level topic proportions describe quantitatively what the documents with molecule *i* talk about. In other words, they help measure the molecule’s chemical and industrial characteristics.

36. This problem does not arise for patents since the data provides the *application* date of patents.

37. Online Appendix Figure B2 shows that coherence increases with the number of topics up to a certain point.

38. I train the algorithm on the set of documents that mention at least one molecule (either a CFC or CFC substitute or HAPs).

39. Catalysts are molecules (typically “metals”, word 7 in topic 1) used to start and maintain chemical reactions. They are often made more effective by being attached to a polymer substrate (which will contain some “carbon” atoms). Hence, the production of such catalysts often involves a “polymerization” process (word 6 in topic 1) where free “radicals” (word 13 in topic 1) ensure the addition of new monomers to the polymer chain.

40. Online Appendix Figure B1 summarizes these steps.

4.4 Similarity between HAPs and CFC Substitutes

This section provides more evidence of the similarity between HAPs and CFC substitutes. First, I document that patents about CFC substitutes and patents about HAPs are associated with similar technology codes. Figure 2 illustrates this point by showing the distribution of 3-digit International Patent Classification (IPC) codes. We note, for example, that both corpora contain many patents related to class C08 (Organic Macromolecular Compounds) and C07 (Organic Chemistry).


Topic proportions provide another, more continuous, way to characterize how similar molecules are. Figure 2b illustrates the variation in topic proportions across molecules. Some HAPs have values of topic proportions that stand out as outliers, indicating that they present semantic contexts that are very different from those of CFC substitutes. In short, topic proportions will be useful to control for differences across molecules.

To quantitatively document the distance between CFC substitutes and HAPs, I compute cosine similarities in the vector space of patent codes and topic proportions for the HAPs included in the difference-in-differences control set and the synthetic control donor pool. Since patent codes provide an easy correspondence to industry codes, I also compute cosine similarities in the vector space of NACE codes. For simplicity, I group the 14 CFC substitutes together and calculate the distance between each HAP and this “Aggregate” CFC substitute.

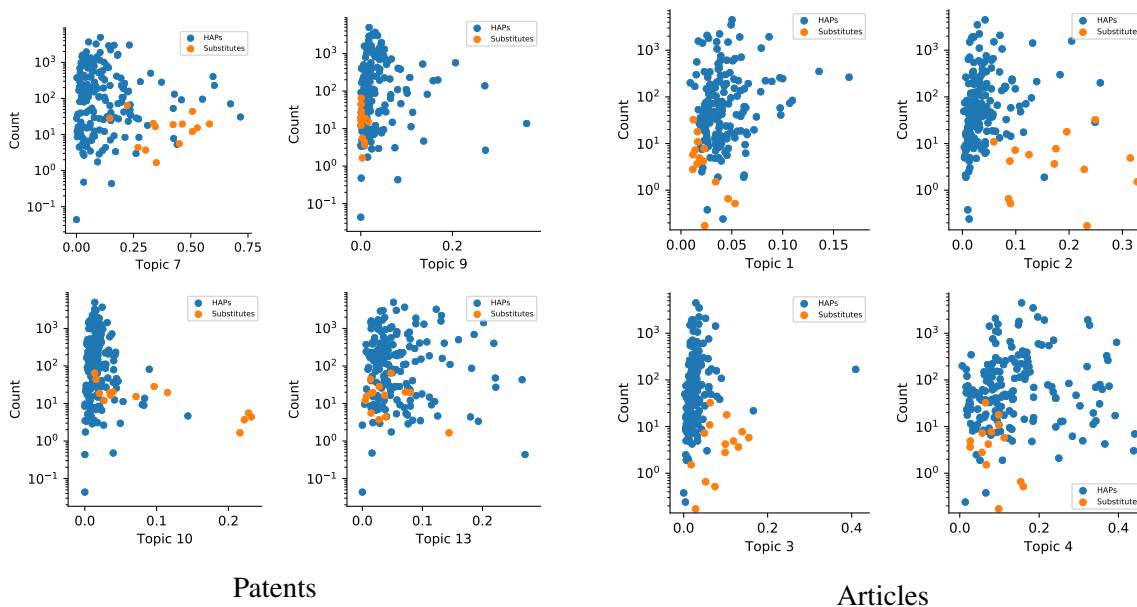
I find that many HAPs have high similarities. For the subset in the DiD control group, the 75th percentile HAP has similarities of 0.47, 0.62, and 0.97 in the space of 4-digit IPC codes, weighted topic proportions, and NACE codes, respectively. For HAPs in the SCM donor pool, the 75th percentile similarities equal 0.73, 0.82, and 0.99 for 4-digit IPC codes, weighted topic proportions, and NACE codes, respectively.⁴¹ Not all HAPs present a high similarity, but my empirical strategy will ensure that more weight is placed on those that do to estimate treatment effects.⁴²

41. Full results are available in Online Appendix Table D1, D2 and D3.

42. As explained in Section 5, the DiD control group is constructed such that it contains the HAPs with the closest pre-trends and the SCM donor pool such that it contains the HAPs with most similar pre-period counts. As a result, similarity in terms of topics or codes is not a criterion for inclusion. Instead, topic proportions are used as control variables in the DiD and in the SCM.



(a) Patent IPC Codes



(b) Topics Proportions and Counts for Selected Topics

FIGURE 2
CFC substitutes vs. HAPs

Note: Figure 2 shows that patents mentioning CFC substitutes and HAPs fall into similar patent codes. HAPs are a group of 171 molecules that have no relationship to ozone and that are used for diverse industrial applications. The figure confirms that the two groups of molecules present remarkable 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 (IPC). Figure 2b illustrates the variation across HAPs and CFC substitutes in topic proportions and document counts. Topic modeling allows to measure the proportions of different topics in documents. They are averaged over all documents mentioning a given molecule (across all years in the sample). The x-axis on the figure above plots the average topic proportion for a given molecule. The y-axis shows the average number of documents mentioning a given molecule. The graphs highlight that CFC substitutes and HAPs don't always have comparable topic proportions indicating that their semantic contexts can be very different. Using topic proportions in the DiD and synthetic control methods therefore provides a way of controlling for variation in how different molecules are mentioned in the text of documents.

5 RESULTS

5.1 Difference-in-Differences

HAPs are related to industrial activities similar to those of CFC substitutes, but some are more similar to CFC substitutes than others. In particular, some have much higher counts than the average CFC substitute.⁴³ In what follows, I exclude from the control group HAPs that have average pre-period counts larger than ten times that of the average CFC substitute. I then rank the remaining HAPs according to how close their pre-trend is to that of the average CFC substitute. The control group is constructed such that it contains the 28 HAPs with the closest slope. I choose the number of HAPs to be 28 so that the control group has twice as many units as the treated group. Figure 3a shows that pre-trends for the treated and control groups are parallel.

I estimate the following two DiD models: a mean shift specification (Equation 3) and a trend-break specification (Equation 4):

$$Count_{mt} = \alpha + \beta_0 \cdot D_m \cdot Post_t + \lambda_t + \lambda_m + \gamma_t \cdot X_{mt} + \varepsilon_{mt} \quad (3)$$

$$Count_{mt} = \alpha + \beta_1 \cdot Years \cdot Post_t \cdot D_m + \beta_2 \cdot Years \cdot Post_t + \beta_3 \cdot Years + \lambda_t + \lambda_m + \gamma_t \cdot X_{mt} + \varepsilon_{mt} \quad (4)$$

$Count_{mt}$ stands for the number of documents mentioning molecule m in year t ; $Post_t$ equals one when $t > 1987$; D_m equals one if the molecule belongs to the treated group; $Years$ is a continuous variable indicating the number of years relative to 1987; λ_m are molecule fixed effects; λ_t are year fixed effects; X_{mt} is a vector of topic proportions. β_0 identifies the DiD estimate. The primary hypothesis is that β_0 and β_1 are positive. Furthermore, Online Appendix Table E1 displays balance tables for topic proportions and highlights that, for most topics, proportions are significantly different across the two groups. They may, therefore, capture relevant variation.

The dependent variable is best suited to be modeled as count data, but since results do not differ much from the simpler easy-to-interpret linear specification with counts in levels,

43. See Online Appendix Figure D1 and D2 for details.

I report the latter as the main specification for ease of interpretation. Results using Poisson pseudo-maximum likelihood regressions are shown in Online Appendix Table E4.

Table 2 displays the main results. Model 1 and 5 correspond to the differences-in-differences specification without any controls. The binary variable “Post 1987 x Substitutes” equals 1 for observations belonging to CFC substitutes after 1987. Model 2 and 6 control for topic proportions and is the preferred specification. The coefficient is large and statistically significant for patents: it corresponds to an additional 21 patents on average per year per substitute from 1988 to 1992. Since there are 14 CFC substitutes in the sample, this implies 294 additional patents a year for CFC substitutes in aggregate. It is equivalent to almost a 400% increase relative to the pre-period mean number of patents (which equals 5.4).

The coefficient for articles is also large and significant: it corresponds to an additional 13 articles on average per year per substitute from 1988 to 1995. Since there are 14 CFC substitutes in the sample, this implies 182 additional patents a year for CFC substitutes in aggregate. It is equivalent to a 576% increase relative to the pre-period mean number of patents (which equals 2.19). Model 3, 4, 7, and 8 present trend-break specifications and confirm that the data can be modeled as a trend break too. The number of patents mentioning CFC substitutes increases with the years after 1987 by 5.8 patents more than the control group. Similarly, the number of articles mentioning CFC substitutes increases with the years after 1987 by two articles more than the control group.

Figure 3b displays the year-by-year coefficients corresponding to Model 2. We note that, for patents, the treatment effect is statistically significant, starting in 1989 and keeps increasing with time. We should expect a delay between the moment firms and inventors decide to redirect their efforts towards CFC substitutes and the moment when they are ready to apply for a patent. However, the time required to obtain any technology worth patenting can vary broadly from technology to technology, even within the same technological field. We can expect some patents to be “low-hanging fruits,” i.e., inventions requiring just a few months of R&D work before applying for a patent. Hence, it is not surprising to observe a significant treatment effect as early as two years after Montreal. Others have documented similarly fast treatment effects. In the context of energy patenting, Popp (2002) estimates that the mean lag between a change

TABLE 1
First Differences Results

	(1) Patents	(2) Patents	(3) Articles	(4) Articles
Post 1987	29.51 (2.11)	6.10 (2.63)	13.02 (1.07)	2.11 (1.58)
Post 1987 x Years		3.95 (0.44)		1.44 (0.28)
Years		-0.03 (0.25)		0.16 (0.06)
Molecule FEs	Yes	Yes	Yes	Yes
Bootstrapped SE	Yes	Yes	Yes	Yes
R-squared	0.64	0.74	0.58	0.63
Observations	322	322	406	406

Standard errors in parentheses. Variable 'Years' is relative to 1987.

Time period: 1976-1998 for patents; 1970-1998 for articles

Note: The table presents regression results for first-difference specifications. Model 1 and 3 confirm that there is a significant and positive mean shift after 1987 in the number of patents and articles mentioning CFC substitutes. Model 2 and 4 indicate 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 articles.


TABLE 2
Difference-in-Differences Results

	(1) Patents	(2) Patents	(3) Patents	(4) Patents	(5) Patents	(6) Articles	(7) Articles	(8) Articles	(9) Articles	(10) Articles
Post 1987 x Substitutes	16.54 (1.74)	21.38 (2.15)	21.12 (2.06)	0.46 (2.68)	3.44 (2.49)	7.58 (1.12)	12.20 (1.53)	12.63 (1.69)	-0.27 (1.63)	1.44 (2.48)
Post 1987 x Substitutes x Years				5.23 (1.01)	5.77 (0.83)				1.82 (0.38)	2.06 (0.50)
Substitutes x Years				0.05 (0.12)	-0.00 (0.13)				-0.03 (0.08)	0.11 (0.16)
Years				-0.07 (0.06)	-0.09 (0.07)				0.18 (0.03)	0.21 (0.04)
Post 1987				1.18 (0.65)	1.21 (0.76)				0.97 (0.45)	0.91 (0.51)
Year FEs	Yes	Yes	Yes	No	No	Yes	Yes	Yes	No	No
Molecule FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Topics (weighted)	No	No	Yes	No	Yes	No	No	Yes	No	Yes
Bootstrapped SE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.84	0.85	0.86	0.85	0.87	0.63	0.63	0.64	0.63	0.64
Observations	714	595	595	714	595	1092	846	846	1092	846


Standard errors in parentheses. Variable 'Years' is relative to 1987.

Time period: 1976-1992 for patents; 1970-1995 for articles

Note: The table presents OLS regression results for difference-in-difference specifications. Model 1 and 5 confirm that there is a significant and positive increase after 1987 in the number of patents and articles mentioning CFC substitutes compared to the control group. Columns 3 and 8 control for topic proportions. This helps ensure that the change in document counts is not driven by a change in demand for particular industrial applications as captured by topic proportions. Topic proportions are missing for years when a molecule is not mentioned in any patent or article. As a result, the numbers of observations are smaller in Columns 3 and 8 compared to 1 and 6. For comparability, Columns 2 and 7 display the same specification as Columns 1 and 6 but using the sample of observations used in the estimation of Columns 3 and 8.



(a) Trends in Counts of Documents Mentioning CFC Substitutes and HAPs



(b) Difference-in-Differences Treatment Effects by Year Controlling for Topic Proportions

FIGURE 3
Difference-in-Differences

Note: Figure 3a display the trends for the treated group (CFC substitutes) and the control group constructed using a subset of the HAP molecules that have counts and trends before 1987 closest to the average CFC substitutes. A version of this figure with trends for CFC compounds is shown in Online Appendix Figure I4. Figure 3b shows the difference-in-differences treatment effects by year controlling for topic proportions. For patents, the treatment effect is statistically significant, starting in 1989 and keeps increasing with time. For articles, the treatment effect is significantly different from zero from 1990 onwards.

in energy prices and patenting on energy-saving technologies occurs in 3.71 years.

For articles, the treatment effect is significantly different from zero from 1990 onwards (see Panel b in Figure 3b). Since the data only contains the publication date of articles (as opposed to submission date), the lag between treatment and its effect also accounts for processing and reviewing time at academic journals which, in chemistry and engineering, takes around nine months on average (Björk and Solomon 2013).

Using HAPs to construct a counterfactual, the treatment effect captures the overall impact of the protocol's signature, the country-by-country ratifications,⁴⁴ and the amendments to the protocol that happened in 1990 and 1992 (in London and Copenhagen, respectively). It also includes, therefore, the implementation of these international treaties into domestic regulations. In the US, this was done through the Clean Air Act amendment of 1990. The counterfactual represents a world without any of those interventions, that is, a world with limited unilateral actions. This serves as a suitable reference point, considering the inherent difficulties associated with addressing global public goods problems. The absence of costly unilateral actions before 1987 is further testimony to those challenges.⁴⁵

5.2 Robustness Checks

I run additional DiD specifications controlling for lags of counts. A typical model of science and innovation is one with positive knowledge externalities: patents or articles lead to more patents and articles as scientists and inventors build on previous work. Controlling for lagged count allows capturing such a cumulative mechanism. Model 2 in Online Appendix Table E2 confirms that the treatment variable remains statistically significant in both patents and articles. Model 3 shows that results are robust to considering, as an alternative outcome variable, counts weighted by the number of forward citations each document receives. Citations can be interpreted as the publication's degree of influence (or "quality"). Panel a and c in Online Appendix Figure E1 show that pre-trends in citation weighted counts look similar across CFC substitutes

44. The USA ratified in April 1988, European countries in December 1988. Online Appendix Table A3 provides the ratification dates for key signatories. It is worth noting that there is little variation over time, as all high-income countries ratified between April and December of 1988.

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

and HAPs for both patents and articles.

Similarly, Model 4 shows that the results are robust to considering patent and article counts weighted by the number of times the molecule appeared in each document (molecule “occurrence”). Panel b and d in Online Appendix Figure E1 show that pre-trends in occurrence-weighted counts look similar across CFC substitutes and HAPs in patents. For articles, a small pre-trend indicates that when articles contain the name of CFC substitutes, they tend to mention these substitutes more often over the years. I provide more details about trends in counts for different thresholds of occurrences in Online Appendix Figure E2.

5.3 Synthetic Control Method

The DiD strategy implemented above assumes that the counts of patents and articles for each molecule are independent. However, molecules are often mentioned together in the same documents: in fact, 40% of patents mention more than one molecule.⁴⁶ Instead of considering molecules as separate units of observations, an alternative strategy is to group them together and count the number of documents mentioning any of the 14 CFC substitutes. This strategy is equivalent to considering the 14 molecules as one treated unit, which I refer to as the “aggregate CFC substitute”. To illustrate, Figure 4 plots both the number of patents mentioning each CFC substitute, as well as the number of patents mentioning any of the 14 CFC substitutes, i.e., the “aggregate CFC substitute”.

I obtain an estimate of the treatment effect on the “aggregate CFC substitute” by implementing a synthetic control method (SCM). This allows me to construct a counterfactual molecule that mimics the evolution of CFC substitutes in aggregate. The synthetic control method emerged as a way to evaluate the effects of interventions that affect aggregate quantities (Abadie, Diamond, and Hainmueller 2010, 2015; Abadie and Gardeazabal 2003; Athey and Imbens 2017). Many interventions are implemented at an aggregate level and impact 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.

46. Out of 3270 patents mentioning CFC substitutes, 1234 mention more than one CFC substitutes. The DiD considers 5999 observations when, in reality, there are only 3270. For articles, out of 998, 226 mention more than one CFC substitutes. Hence the DiD considers 1266 observations when, in reality, there are only 998.


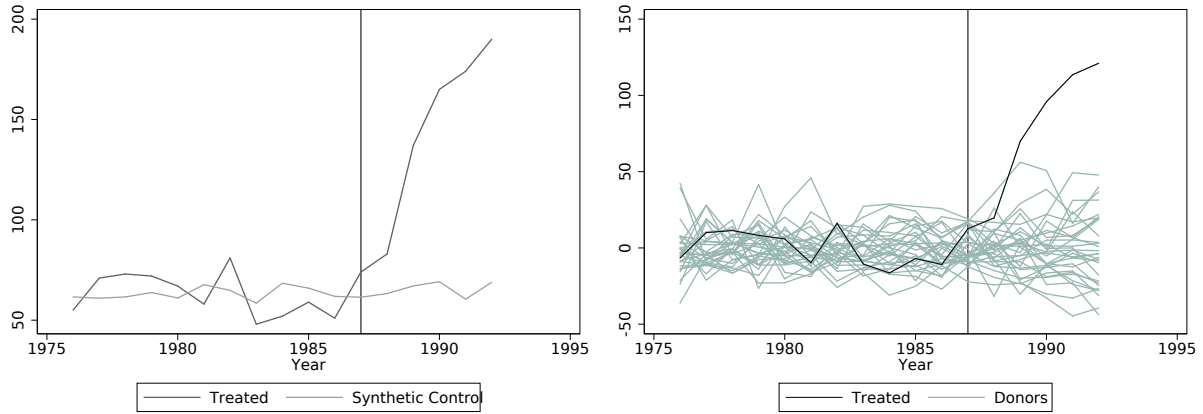



FIGURE 4

Patent Counts for Each CFC Substitute and for the “Aggregate” CFC Substitute

Note: The graph illustrates the difference between considering the 14 separately or together combined as one treated molecule. Since the names of different CFC substitutes often appear simultaneously in the same documents, the time series of CFC substitutes are not independent of each other. The thick line labeled “Any CFC substitutes” 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.” Online Appendix Figure F1 displays a similar graph for articles.



(a) Patents: Raw Effect (left) and Placebo Tests (right)



(b) Articles: Raw Effect (left) and Placebo Tests (right)

FIGURE 5

Synthetic Control Method Graphs for CFC Substitutes

Note: 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 of 1989. This indicates that similarly, as in the DiD, the treatment effect is statistically significant, starting in 1989. For articles, the treatment effect is statistically significant, starting in 1992.

The synthetic control method consists of 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 in the pre-treatment years along some selected covariates and the outcome variable.⁴⁷ The outcome variable here is the number of patents (or articles) that mention any of the 14 molecules. The synthetic control is constructed by fitting the values of pre-treatment counts and topic proportions. 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 only up to 1985 to fit the synthetic control.⁴⁸ Topic proportions are averaged over the entire pre-1985 period, while the outcome, count, is not.⁴⁹

As explained by Abadie, Diamond, and Hainmueller (2015), reducing the size of the donor pool can limit the risk of over-fitting and interpolation biases. Following their advice, I use a smaller donor pool containing only the thirty HAPs closest to the treated unit in terms of counts.⁵⁰ For inference, I follow the non-parametric approach method suggested by Abadie, Diamond, and Hainmueller (2010) and Abadie, Diamond, and Hainmueller (2015). The approach is akin to implementing placebo tests wherein each unit in the control group is assumed to have received the treatment in the year 1987. The “true” treatment effect can then be compared to the distribution of placebo treatment effects. 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, Diamond, and Hainmueller (2010), I compute the ratios of post-RMSPE over pre-RMSPE and examine where the treated unit lies in the distribution of those ratios.⁵¹

Figure 5 graphically displays the results of the synthetic control method for CFC substitutes. The graphs on the left-hand side represent the raw effect, which is the observed time series of

47. Online Appendix F provides more details on the theoretical foundations for the synthetic control method.

48. As a robustness check, I used only data from the first part of the pre-period (from 1976 to 1980 for patents and from 1970 to 1978 for articles) to fit the synthetic control. In both cases, the synthetic control predicted the trend well in the second part of the pre-period, and the results were virtually unchanged. This analysis is shown in Online Appendix Figure F5.

49. Online Appendix Table F3 displays the value of each variable’s contribution to the synthetic control.

50. I also check that there is no risk of extrapolation. See Online Appendix Table F1

51. 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 pre-treatment periods. A post-RMSPE can be similarly defined for periods going from $T_0 + 1$ to the end of time-series available.

the treated group along with the time series of the constructed synthetic control. Graphs on the right-hand side show the placebo tests to evaluate the significance of the results; the black lines represent 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 treatment effect on CFC substitutes appears significant for both patents and articles. We note that the black line rises above most other lines, after 1989. This indicates that, as in the DiD, the treatment effect is statistically significant, starting in 1989. For articles, the treatment effect is statistically significant, starting in 1992.


Graphically, the treatment effect corresponds to the area between the two curves on the left-hand side graphs. Numerically, it corresponds to 84 (resp. 49) additional patents (resp. articles) per year on average between 1988 and 1992 (resp. 1995). This roughly corresponds to 135% and 177% increases in patents and articles. To examine more closely the constructed control unit, I collect information about the top HAPs that enter its composition (see Online Appendix Table F2). We note that the synthetic control picked up HAPs that have broad industrial applications (not unlike CFC substitutes). In particular, we find industrial applications similar to CFC substitutes (e.g., coatings or solvents).

6 MECHANISMS AND ALTERNATIVE EXPLANATIONS


6.1 Mechanisms

This section provides descriptive details about the increase in research and innovation fostered by Montreal. In Figure 6, we could already note that the increase in patenting applies to almost all CFC substitutes. Three molecules see strong increases in the three years that follow the signature of Montreal: HFC 134a, HCFC 141b, and HCFC 123.⁵² These substitutes were the most promising in terms of thermodynamic properties. Their vapor pressures were similar to CFC-11 and CFC-12, which implied they could be used with little change to equipment and operations (Manzer 1990; Parson 2003, p.54). Despite their promise, however, firms had no


52. HFC 134a patents increase from an average of one patent a year before 1987 to about 30 per year during the period 1987 to 1990. Similarly, HCFC 141b patents increase from 7 to 30, and HCFC 123 patents from 2.6 to 22.5.




(a) New Entrants Over Time




(b) Patenting by Type of Firms




(c) Top Patenting Firms



(d) Patenting on CFCs and CFC Substitutes



(e) Monthly Counts - All Assignees



(f) Monthly Counts: Old vs. New Entrants

FIGURE 6

Firms and Patenting on CFC Substitutes

Note: Figure 6a displays the number of assignees that are “new”, i.e., assignees that apply for a patent on CFC substitutes or HAPs for the very first time in a given year. The figure shows that, after 1987, many firms with no prior experience on CFC substitutes begin patenting. The data for HAPs is normalized such that y-axis values are equal to those of CFC substitutes in 1976. To limit noise, the sample used to generate the table contains only documents with at least three occurrences of CFC substitutes. Figure 6b highlights that those new entrants are responsible for 43% of CFC substitutes patents after 1987 (top right flow colored in yellow) while firms that had already patented on CFC substitutes are responsible for 57% of those patents. Figure 6c shows the patenting activities of the top 10 firms over time. The scatterplot on Figure 6d indicates that firms with more CFC patents are also those with more CFC substitutes patents. For this graph, the patent counts include patents filed in any year. Figure 6e and 6f show the monthly trends in count of patents mentioning CFC substitutes. Panel 6f shows the monthly count of patents mentioning CFC substitutes for firms that patented on CFC substitutes before 1987 vs. those who did not. The period “Before 1987” includes the year 1987. To limit noise, the sample used to generate the table contains only documents with at least three occurrences of CFC substitutes.

prior commercial experience with those compounds, and deploying them required research and development on all aspects, from toxicity, process testing, and synthesis routes.

Similar technological codes are associated with patents before and after 1987, although many codes with low or no frequency before 1987 become important after 1987. That is the case, for example, of code C10M which corresponds to “Lubricating compositions; Use of chemical substances either alone or as lubricating ingredients in a lubricating composition.”⁵³ I also find that the increase in the number of patents mentioning CFCs applies to all countries and a substantial increase for patents with assignees located in Japan and the UK.⁵⁴

Next, I examine the assignees behind CFC substitutes patenting. The first takeaway is that almost all patents come from private sector firms, both before and after Montreal (97% and 96%, respectively). Interestingly, Figure 6a reveals a post-1987 significant increase of “new entrants”, i.e., firms that did not file any patents for CFC substitutes prior to 1987. In total, there are 566 firms identified as the assignees of post-1987 CFC substitutes patents. 87% of these firms are “new entrants” who have not filed CFC substitutes patents before 1987, while a minority of 13% are “incumbents” who have filed such patents previously. Figure 6b provides a breakdown of the post-1987 CFC substitutes patents by type of firms, showing that a majority of these patents (57%) are filed by incumbents (bottom right orange flow), while new entrants are responsible for 43% (top right yellow flow). Figure 6c zooms in on the top 10 patenting firms and shows how their patenting evolved over time. We note that Dupont and Allied are by far the most active and the quickest to patent after Montreal, but others such as Dow, Elf Atochem, or Daikin also ramp up patenting quickly. Finally, the scatterplot in Figure 6d indicates that firms with more CFC patents are also those with more CFC substitutes patents. I also find that, at the firm-level, the number of patents before 1987 (either on CFCs or on CFC substitutes) correlates positively with the number of CFC substitutes patents after 1987.

53. Other examples include “Macromolecular compounds obtained otherwise than by reactions only involving carbon-to-carbon unsaturated bonds” (C08G), “Cleaning or de-greasing of metallic material by chemical methods other than electrolysis” (C23G), and “Detergent compositions; Use of single substances as detergents; Soap or soap-making; Resin soaps; Recovery of glycerol” (C11D).

54. See Online Appendix Table G5 and Online Appendix Figure G1 and G2 for more details.

6.2 Were CFC Substitutes Kept “Secret”?

In this section, I examine whether firms may have incurred significant innovation efforts related CFC substitutes before the signature of the Montreal Protocol, without patenting but instead keeping their technologies as trade secrets. At the end of the 1970s, a few firms announced R&D investments into CFC substitutes. Although the same firms, soon after, announced the termination of those same R&D programs, it has been suggested that they developed key technologies which they kept secret. Here, I argue that, if that had been the case, we should expect a one-time increase in patent counts in the immediate aftermaths of Montreal.

As explained by Parson (2003), trade secrets are difficult to keep when developing CFC substitutes development. The key remaining barriers were to prove suitability for specific applications. This could not be done in secret as it required partnerships with customers (e.g., electronic manufacturers that used CFCs as a cleaning agent). Developing new synthesis processes could, in theory, be done in secret. However, several firms were working on the same molecules. When competitors work on closely related projects, delaying patenting increases the risk that a competitor patents first. Incentives to be first to patent were therefore particularly strong. If, until Montreal, firms thought policy pressure was low, they may have elected not to patent. However, once the protocol is signed, they have tangible incentives to patent any old technologies that they may have previously been developed as fast as possible to outrun possible competitors.

If CFC substitutes were kept secret, we should then observe an immediate peak in the number of patents in the few months following the signature of the agreement. Figures 6e and 6f plot the number of patents mentioning CFC substitutes month by month in the two years that followed Montreal. On the first graph, we note the absence of a patenting peak after 1987, which provides support against the hypothesis that technologies related to CFC substitutes had been kept secret. The second graph presents trends for assignees that never patented on CFC substitutes before 1987 and those who did. Suppose the R&D carried out before Montreal was a key driver to the post-Montreal increase in patenting. In that case, we may observe significant differences in patenting trends between firms with and firms without prior patenting experience on CFC substitutes. Although a gap seems to build up over time, trends look mostly similar.

While dozens of firms actively patented on CFC substitutes, two firms, in particular, dominate the landscape: DuPont and Allied. Online Appendix Figure G3 zooms in on these two actors. Figure G3a shows that most patents granted to DuPont and Allied were applied for after 1989, and in particular, Figure G3b shows that there is no sudden peak patenting right after Montreal. Instead, we observe a gradual ramping up of patenting activity. Figure G3c illustrates that the patents from DuPont and Allied, which received the highest number of citations, mostly originate from 1989 to 1991. Figure G3d 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 Allied likely had a first-mover advantage on some technologies.⁵⁵ However, the ramping up in patenting activity from 1990 onwards supports the claim that, even for DuPont and Allied, most of the innovative activity started after Montreal. Of course, this does not mean that no knowledge existed at all. After all, some patents were filed before 1987. But it indicates that it is unlikely that particular pieces of knowledge or technologies had been hidden.

6.3 Was Consumer Pressure a Potential Driver?

The science of ozone made much progress during the 1980s. In particular, in 1985, scientists detected an extensive depletion of ozone over Antarctica (the “hole”), and importantly, they were able to causally attribute it to CFCs in March 1988 (the “discovery”). The image of the Earth seen from space with a massive hole (artificially colored in blue for the occasion) became world-famous and moved public opinion. The perceived benefits of phasing out CFCs certainly increased and made the issue more salient in the public’s eyes. To what extent, then, did consumer pressure drive innovation in the aftermaths of Montreal?

There are very few empirical analyses suggesting that consumer pressure is effective (Lyon and Maxwell 2002; Popp, Hafner, and Johnstone 2011), and they all deal with local pollutants

55. 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 and Vietor 1989). It is therefore unlikely that financial markets would capture much impact. In addition, it would be 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).

such as toxic chemical emissions. We could reasonably expect consumer pressure to be less effective for a 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. To investigate the role of consumer pressure for innovation on CFC substitutes, I leverage the fact that not all molecules were involved in consumer-facing applications (e.g. industrial applications). Restricting the analysis to substitutes whose applications were not exposed to consumers allows setting aside the possible influence of public opinion.

I use data available on the EPA SNAPs website to identify which CFC substitutes were not exposed to consumers. The EPA website lists a total of 1001 requests corresponding to a given substance for a given application. As part of the implementation of the Montreal Protocol, the EPA was in charge of controlling which substitutes firms could use. To this end, any firm using or producing substitutes had to request authorization for specific substance-application pairs. An example is HCFC-225cb for Electronics Cleaning. I manually classify applications as either consumer exposed, not consumer exposed, or undetermined. In doing so, I identify 6 CFC substitutes that are not consumer exposed. The typical applications requested for these CFC substitutes are related to industrial activities invisible to consumers: e.g., centrifugal chillers, foam blowing agents for rigid polyurethane, or precision cleaning.

I find that, for almost all CFC substitutes that were arguably not exposed to consumers, the number of patents increases sharply after 1987. This is indicative that consumer pressure and public opinion did not play an essential role in driving innovation in the aftermath of the Montreal protocol.⁵⁶

In addition, I also investigate whether consumer exposure played a role in the induced innovation response by examining whether CFC substitutes that were exposed to consumers had a stronger response than those that were not. To do this, I estimate specifications similar to those in Columns 1 and 2 of Table 2, but with an additional interaction term with a binary variable that equals 1 for CFC substitutes with applications exposed to consumers and 0 for those that are not.

56. The trends for both exposed and non-exposed CFC Substitutes are shown in Online Appendix Figure G16.

The coefficient on the interaction term is not interpreted as causal here because there is no quasi-random variation in the assignment of consumer exposure status. However, the analysis shows that CFC substitutes exposed to consumers have a stronger response than those not exposed, but the effect disappears when controlling for topic proportions, indicating that other molecule-level characteristics, such as the types of industrial applications or the thermodynamic or safety profiles, may explain why consumer-exposed molecules exhibit a stronger response.⁵⁷

7 TOWARDS A THEORY OF ENVIRONMENTAL AGREEMENTS WITH INDUCED INNOVATION

The role of the Montreal Protocol in solving the crisis has been intensely discussed. In particular, in a seminal article, Barrett (1994) developed a theory of international environmental agreements which interpreted Montreal's success as a case of cheap cooperation. Murdoch and Sandler (1997) also argued that some countries, and especially the USA, accepted to sign Montreal because the costs of doing so were low. Core to their argument is the existence of a cost-benefit analysis issued by the EPA in 1987, concluding that the estimated benefits (mostly from cancers avoided) overwhelmed the estimated costs to the industry. Montreal being "cheap" has often been rephrased as the idea that substitutes existed. However, it is more exact to interpret it as an agreement that industries and governments expected to be technologically feasible within the time frame agreed upon.

This paper highlights that, even if Montreal encoded what some thought feasible from a technological point of view, much experimentation and R&D investments were still required. As a result, the agreement induced innovation. By considering the effects of international agreements on innovation, we are given the opportunity to re-interpret Montreal's success. The first key point is that Montreal's original targets can be seen as modest. Diplomats, in fact, failed to negotiate a full ban in 1987 (against which the industry actively lobbied). Real success instead should be seen in the later amendments (London in 1990 and Copenhagen

57. Full results are shown in Online Appendix Table G6.

in 1992), where more ambitious reduction targets were agreed and other molecules added to the list of regulated compounds. Montreal may not have been very ambitious, but, it set up a credible enforcement mechanism that changed firms' expectations about CFCs and, hence, induced innovation. In doing so, it contributed to lowering down the perceived costs associated with a complete phase-out.

Conceptualizing agreements as a vehicle for inducing innovation bears important implications for interpreting the theory of international agreements. Prior work has studied many aspects of cooperation, including how technology shapes cooperation outcomes (Harstad, Lancia, and Russo 2019). But the role of innovation is usually not explicitly considered. Induced innovation suggests that as efforts are made to reduce emissions, the costs associated with those efforts will decrease over time. In other words, abatement costs should decrease over time as abatement actions are undertaken. But previous studies have mainly represented the costs of emission reduction as remaining constant over time.

In this section, I revisit some of these prior studies to discuss how induced innovation would affect models of international environmental agreements. I first start with the seminal contribution of Barrett (1994). In this simple setup, N countries pollute and decide whether or not to pay for pollution abatement. The costs are incurred by each country separately, while the benefits of abatement accrue to all. Free-riding incentives, therefore, arise: countries would be better off if all were to abate a high amount of emissions (the cooperative level), but the Nash equilibrium of the game leaves all countries at a lower amount of abatement (the non-cooperative level).

Barrett (1994) then examines the conditions for successful self-enforcing agreements. He shows that large coalitions occur only when the cost of abatement is low; unfortunately, this is also when countries would find it beneficial to reduce emissions unilaterally. In other words, there is little difference between the non-cooperative and cooperative equilibria: the gains from cooperation are low. This is the essence of the dismal conclusion about self-enforcing agreements: they only appear when we don't really need them. And in this context, the Montreal Protocol has been featured as a poster example.

This paper instead emphasizes that success in the fight against ozone depletion should be

seen in the London and Copenhagen amendments, which took place in 1990 and 1992, respectively, where targets were renegotiated. Looking back from 1987, these more ambitious targets were not considered economically viable and, in line with theoretical predictions, were not incorporated into the Montreal Protocol.

Adding induced innovation to Barrett (1994)'s model demonstrates that, despite starting with agreements where cooperation gains are small, over time, induced innovation allows us to successfully negotiate targets that used to be unattainable. A simple way to model induced innovation is to assume that the marginal costs of abatement decrease with the amount of abatement done in the previous periods.

By enforcing some emission reductions in period t (even if only the “cheap” emissions), the agreement forces firms to do new things, experiment, and develop new or improve old technologies. These processes pave the way for lowering the marginal cost of abatement in the next period, and large coalitions can arise where they used to be unlikely. In other words, in 1987, the London and Copenhagen targets were too expensive, but induced innovation made them within reach of an agreement a few years later.

Other theoretical work has examined how technology shapes outcomes in international environmental agreements. Still, these papers do not explicitly consider “innovation”, or in other words, the possibility that abatement costs may decrease over time. For example, Harstad, Lancia, and Russo (2019) study how investments in various types of technologies influence whether the agreement is self-enforcing. They show that the best subgame-perfect equilibrium requires countries to strategically overinvest in “clean” technologies as it reduces the temptation to defect.

In this context, investing in technologies is akin to building energy infrastructure before the emission decision is taken. The comparative statics highlight that it is harder to motivate compliance when investment costs in clean technologies are high; this result is equivalent to the findings in Barrett (1994) that there are no large self-enforcing coalitions when abatement costs are high. In Harstad, Lancia, and Russo (2019), investment costs are fixed, but induced innovation could simply be modeled as such costs going down over time (e.g., by making them a function of prior investments). This, in turn, would make it easier to motivate compliance.

To summarize, the success story of Montreal is better summarized as a repeated cooperation game where induced innovation enabled emission reduction ambitions to ratchet up. As such, the evidence presented points to co-evolution of international environmental agreements and innovation efforts. At each stage, binding reductions forced firms to innovate and develop technologies to comply. With each realized innovative step, expectations for further progress increased, and the expected costs of abatement decreased. More aggressive reductions were then perceived as affordable, and stakeholders (both public and private) became willing to bind themselves to more ambitious targets.

8 CONCLUSION

Tackling environmental problems often relies on developing and diffusing new technologies. Understanding the drivers of technological change is, therefore, a critical input to improving the prospects of environmental cooperation. In this paper, I document that the Montreal Protocol led to the development of CFCs substitutes. This empirical evidence goes against the conventional narrative that alternatives technologies were readily available before the treaty. By showing that a low-ambition but binding agreement such as the Montreal Protocol did encourage the development of technological solutions, this paper suggests such agreements are potent tools that dynamically increase the net benefits of environmental protection and may therefore also be useful to deal with current problems such as climate change.

SUPPLEMENTARY MATERIAL

The Online Appendix for this article can be found at eugeniedugoua.com/papers/Dugoua_Innovation_Montreal_SOM.pdf.

REFERENCES

- Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2010. "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program." *Journal of the American Statistical Association* 105 (490): 493–505.
- . 2015. "Comparative Politics and the Synthetic Control Method." *American Journal of Political Science* 59 (2): 495–510.
- Abadie, Alberto, and Javier Gardeazabal. 2003. "The Economic Costs of Conflict: A Case Study of the Basque Country." *The American Economic Review* 93 (1): 113–132.
- Acemoglu, Daron. 1998. "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality." *The Quarterly Journal of Economics* 113 (4): 1055–1089.
- Aghion, Philippe, Antoine Dechezleprêtre, David Hemous, Ralf Martin, and John Van Reenen. 2016. "Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry." *The Journal of Political Economy* 124 (1): 1–51.
- Aichele, Rahel, and Gabriel Felbermayr. 2011. "Kyoto and the Carbon Footprint of Nations." *Journal of Environmental Economics and Management* 63 (3): 336–354.
- Ambec, Stefan, Mark A Cohen, Stewart Elgie, and Paul Lanoie. 2013. "The Porter Hypothesis at 20: Can Environmental Regulation Enhance Innovation and Competitiveness?" *Review of Environmental Economics and Policy* 7 (1): 2.
- Andersen, Stephen O, and K Madhava Sarma. 2012. *Protecting the Ozone Layer: The United Nations History*. Edited by Lani Sinclair. Earthscan Publications Ltd.
- Athey, Susan, and Guido W Imbens. 2017. "The State of Applied Econometrics: Causality and Policy Evaluation." *The Journal of Economic Perspectives: A Journal of the American Economic Association* 31, no. 2 (May): 3–32.
- Auffhammer, Maximilian, Bernard J Morzuch, and John K Stranlund. 2005. "Production of chlorofluorocarbons in anticipation of the Montreal Protocol." *Environmental & Resource Economics* 30 (4): 377–391.

- Azoulay, P, J S Graff Zivin, and J Wang. 2010. "Superstar Extinction." *The Quarterly Journal of Economics* 125 (2): 549–589.
- Azoulay, Pierre, Christian Fons-Rosen, and Joshua S Graff Zivin. 2019. "Does Science Advance One Funeral at a Time?" *The American Economic Review*.
- Barrett, Scott. 1994. "Self-enforcing International Environmental Agreements." *Oxford Economic Papers*, 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.
- Battaglini, Marco, and Bård Harstad. 2016. "Participation and Duration of Environmental Agreements." *J. Polit. Econ.* 124, no. 1 (February): 160–204.
- Benedick, Richard Elliot. 2009. *Ozone Diplomacy: New Directions in Safeguarding the Planet*. Harvard University Press.
- Björk, Bo-Christer, and David Solomon. 2013. "The Publishing Delay in Scholarly Peer-reviewed Journals." *Journal of Informetrics* 7 (4): 914–923.
- Blei, David M. 2012. "Probabilistic Topic Models." *Communications of the Acm* 55 (4): 77–84.
- Blei, David M, and John D Lafferty. 2009. "Topic Models." In *Text mining*, 101–124. Taylor / Francis Group.
- Calel, Raphael, and Antoine Dechezleprêtre. 2016. "Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market." *The Review of Economics and Statistics* 98 (1): 173–191.
- Cohen, Wesley M, Richard R Nelson, and John P Walsh. 2000. "Protecting Their Intellectual Assets: Appropriability Conditions and Why Us Manufacturing Firms Patent (or Not)." *Nber Working Paper* 7552.

- Dekker, Thijs, Herman R J Vollebergh, Frans P de Vries, and Cees A Withagen. 2012. “Inciting Protocols.” *Journal of Environmental Economics and Management* 64 (1): 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?” *Journal of Public Economics* 87 (9): 2031–2048.
- Hall, Bronwyn H, and Adam B Jaffe. 2012. “Measuring Science, Technology, and Innovation: A Review.” *Report Prepared for the Panel on Developing Science, Technology, and Innovation Indicators for the Future*.
- Harstad, Bård, Francesco Lancia, and Alessia Russo. 2019. “Compliance Technology and Self-enforcing Agreements.” *Journal of the European Economic Association* 17, no. 1 (February): 1–29.
- Heal, Geoffrey. 2016. *Endangered Economies: How the Neglect of Nature Threatens Our Prosperity*. Columbia University Press.
- Henderson, Rebecca, Adam B Jaffe, and Manuel Trajtenberg. 1998. “Universities As a Source of Commercial Technology: A Detailed Analysis of University Patenting, 1965–1988.” *The Review of Economics and Statistics* 80 (1): 119–127.
- Hicks, John R. 1932. “The Theory of Wages.”
- Iaria, A, C Schwarz, and F Waldinger. 2018. “Frontier Knowledge and Scientific Production: Evidence from the Collapse of International Science.” *The Quarterly Journal of Economics* 133 (2): 927–991.
- Jaffe, Adam B, Richard G Newell, and Robert N Stavins. 2002. “Environmental Policy and Technological Change.” *Environmental & Resource Economics* 22 (1): 41–70.
- Jaffe, Adam B, and Karen Palmer. 1997. “Environmental Regulation and Innovation: A Panel Data Study.” *The Review of Economics and Statistics* 79 (4): 610–619.

- Kellenberg, Derek, and Arik Levinson. 2014. "Waste of Effort? International Environmental Agreements." *Journal of the Association of Environmental and Resource Economists* 1 (1/2): 135–169.
- Lyon, Thomas P, and John W Maxwell. 2002. "'Voluntary' Approaches to Environmental Regulation: A Survey." In *Economic Institutions and Environmental Policy*. Ashgate Publishing.
- Manzer, L E. 1990. "The CFC-Ozone Issue: Progress on the Development of Alternatives to CFCs" [in en]. *Science* 249, no. 4964 (July): 31–35.
- Miller, Alan S, and Irving M Mintzer. 1986. "The Sky Is the Limit." *World Resources Institute, Washington Dc*.
- Milliman, Scott R, and Raymond Prince. 1989. "Firm Incentives to Promote Technological Change in Pollution Control." *Journal of Environmental Economics and Management* 17 (3): 247–265.
- Molina, Mario J, and F Sherwood Rowland. 1974. "Stratospheric Sink for Chlorofluoromethanes: Chlorine Atom-catalysed Destruction of Ozone." *Nature* 249 (28): 810–812.
- Moser, Petra. 2012. "Innovation without Patents: Evidence from World's Fairs." *The Journal of Law and Economics* 55 (1): 43–74.
- Mulder, Karel F. 2005. "Innovation by Disaster: The Ozone Catastrophe As Experiment of Forced Innovation." *Int. J. Environ. Sustain. Dev.* 4 (1): 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." *Journal of Public Economics* 63 (3): 331–349.
- . 2009. "The Voluntary Provision of a Pure Public Good and the Montreal Protocol: Behavioral and Data Concerns." *Oxf. Econ. Pap.* 61 (1): 197–200.
- Parson, Edward A. 2003. *Protecting the Ozone Layer: Science and Strategy*. Oxford University Press.

- Pool, R. 1988. "The Elusive Replacements for CFCs: As CFCs are phased out to protect the ozone layer, chemical companies scramble to find substitutes for these vital compounds and to develop economical ways to make them in large quantities" [in en]. *Science* 242, no. 4879 (November): 666–668.
- Popp, David. 2002. "Induced Innovation and Energy Prices." *The American Economic Review*, 160–180.
- . 2005. "Lessons from Patents: Using Patents to Measure Technological Change in Environmental Models." *Ecological Economics: The Journal of the International Society for Ecological Economics* 54 (2): 209–226.
- . 2019. "Environmental Policy and Innovation: A Decade of Research." *International Review of Environmental and Resource Economics* 13 (3-4): 265–337.
- Popp, David, Tamara Hafner, and Nick Johnstone. 2011. "Environmental Policy Vs. Public Pressure: Innovation and Diffusion of Alternative Bleaching Technologies in the Pulp Industry." *Research Policy* 40 (9): 1253–1268.
- 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*, edited by Bronwyn H Hall and Nathan Rosenberg, 2:873–937. North-Holland, January.
- Porter, Michael E, and Claas Van Der Linde. 1995. "Toward a New Conception of the Environment-competitiveness Relationship." *The Journal of Economic Perspectives: A Journal of the American Economic Association*, 97–118.
- Redner, Sidney. 2005. "Citation Statistics From More Than a Century of Physical Review." *Bulletin of the American Physical Society*.
- Reinhardt, Forest, and Richard H K Vietor. 1989. "Du Pont Freon Products Division (b)." *Harvard Business School Case 389-111*, 261–286.

- Röder, Michael, Andreas Both, and Alexander Hinneburg. 2015. "Exploring the Space of Topic Coherence Measures." In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*, 399–408. New York, NY, USA: ACM.
- Sampat, Bhaven N. 2018. "A Survey of Empirical Evidence on Patents and Innovation." *NBER Working Paper 25383*.
- Shapiro, Michael, and Ellen Warhit. 1983. "Marketable Permits: The Case of Chlorofluorocarbons." *Natural resources journal* 23 (3): 577–591.
- Smith, Brigitte. 1998. "Ethics of Du Pont's CFC Strategy 1975-1995." *J. Bus. Ethics* 17 (5): 557–568.
- Sunstein, Cass R. 2007. "Of Montreal and Kyoto: A Tale of Two Protocols." *Harvard Environmental Law Review*. 31:1.
- 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." *The American Economic Review*, 450–460.
- Wagner, Ulrich J. 2016. "Estimating Strategic Models of International Treaty Formation." *The Review of Economic Studies* 83 (4): 1741–1778.
- Wang, Dashun, Chaoming Song, and Albert-László Barabási. 2013. "Quantifying Long-term Scientific Impact." *Science* 342 (6154): 127–132.
- Williams, Heidi L. 2013. "Intellectual Property Rights and Innovation: Evidence from the Human Genome." *The Journal of Political Economy* 121 (1): 1–27.
- . 2017. "How Do Patents Affect Research Investments?" *Annual Review of Economics* 9:441–469.