Paper to be presented at the

DRUID Society Conference 2014, CBS, Copenhagen, June 16-18

Technology and learning in a long run industry pollution transition: evidence from the global phase-out of ozone depleting substances

David Grover
LSE
Grantham Research Institute
d.grover@lse.ac.uk

Abstract

I estimate a marginal abatement cost (MAC) curve using detailed, project-level data from the global phase-out of ozone depleting substances (ODSs) across 145 low-income countries, including China, Brazil and India. The data cover an estimated 85 percent of the ODSs by weight phased out by these countries. Exploiting a quasi-natural experiment in the implementation of the Montreal Protocol, I find that average project unit abatement cost fell in the initial period of the phase-out before eventually rising, forming a convex curve overall. One explanation for this, which I test, is that pre-existing technology possibilities that could be adapted to phasing out ODSs, flooded into the abatement arena in the early years, pushing down the cost of abatement through a kind of meta-learning. Holding the appearance of these new technology ?possibilities? constant, I also find that each repeat ?application? of each possibility associates with lower abatement cost, as through a learning-by-doing effect. These findings elaborate the theoretical basis for predicting how the cost of abating greenhouse gasses (GHG) may change over time, including why the cost of abatement now is so low.

Jelcodes:O31,Q55

Technology and learning in a long run industry pollution transition: Evidence from the phase-out of ozone depleting substances

Abstract

I estimate a marginal abatement cost (MAC) curve using detailed, project-level data from the global phase-out of ozone depleting substances (ODSs) across 145 low-income countries, including China, Brazil and India. The data cover an estimated 85 percent of the ODSs by weight phased out by these countries. Exploiting a quasi-natural experiment in the implementation of the Montreal Protocol, I find that average project unit abatement cost fell in the initial period of the phase-out before eventually rising, forming a convex curve overall. One explanation for this, which I test, is that pre-existing technology possibilities that could be adapted to phasing out ODSs, flooded into the abatement arena in the early years, pushing down the cost of abatement through a kind of meta-learning. Holding the appearance of these new technology 'possibilities' constant, I also find that each repeat 'application' of each possibility associates with lower abatement cost, as through a learning-by-doing effect. These findings elaborate the theoretical basis for predicting how the cost of abating greenhouse gasses (GHG) may change over time, including why the cost of abatement now is so low.

JEL codes: O31, O33, Q55

Key words: innovation, pollution, technological change, abatement cost

Outline

Tables

Figures

1. Background

The marginal abatement cost (MAC) curve[1] has become influential among policymakers interested in predicting the cost of reducing greenhouse gas (GHG) emissions (Enkvist et al 2007; Pye et al 2008; Stern 2006). A MAC curve typically describes the amount of abatement that a unit of money would buy from some point on the trajectory of abatement achieved (Metcalf and Weisbach 2008). Some renderings group the amount of abatement that could be achieved at a similar unit cost into 'blocks' or 'steps' (IPCC 2013). A MAC curve typically implies that the unit cost of abatement is low over an initial quantity of abatement achieved, rising the closer total abatement achieved gets to 100 percent (Magat 1978; Vijay et al 2010; Popp 2010). MAC (or similar) curves are used to parameterise integrated assessment models (Nordhaus 1991). They influence how policymakers allocate scarce R&D resources. They form the basis of the theoretical relationship between quantity abated and marginal unit abatement cost for all pollutants.

It is not widely appreciated how 'empirical' MAC curves are constructed. Most MAC curves are constructed in one of two ways. I will call these forward-looking and the backward-looking approaches. In the forward-looking approach, estimates are elicited from experts about the quantities of abatement that they believe can be achieved at different unit costs using different abatement measures. The experts are typically engineers. MAC (-like) curves constructed by the forward-looking approach give a misleading sense of certainty about how much is known about what the unit cost of abatement will be at different points in the abatement trajectory. This is because a MAC curve that is constructed from expert opinion can only ever incorporate information that is available at the time the curve is constructed. The curve can never incorporate information about the possibilities for reducing pollution that will be available at the point in time when the abatement that is the subject of the curve will be performed. Forward-looking MAC curves are static. They cannot account for the abatement possibilities that will be present at points in time later than the point in time when the curve is constructed.

In the backwards-looking approach the curve is constructed from information about the cost of abatement activity that already occurred. The empirically observable unit cost of abatement arises from the dynamic interaction among unit cost, the quantity of abatement achieved, and the appearance (or disappearance) of improved or altogether new technologies for reducing pollution. Cost is one of the parameters that arises from the dynamic complexity that integrated assessment

---

[1] Also the marginal cost curve (Nordhaus 1991), the pollution abatement substitution curve (Dellink 2005) and the pollution-saving curve (Oliver et al 1983).

models seek to represent (Carraro et al 2010; Gillingham et al 2008; Grubb et al 2006). Curves constructed by the backwards-looking approach describe a pattern in costs born historically as a result of this complexity rather than costs that are expected to be born in the future.

The problem with backward-looking MAC curves is that the information used to construct them is almost always incomplete, which gives rise to an analytical ‘leakage’ problem. The cost of abatement activity that is observable is often interlinked with, and determined by, the cost of abatement that occurs in an unobservable ‘elsewhere’ (Klepper and Peterson 2006). For example, if the aim is to construct a MAC curve for a pollutant for a single firm, an attractive ‘abatement’ option for the firm will be to outsource the polluting activity to a supplier firm (Copeland and Taylor 2004). This means that the true cost of dealing with a pollutant as it is steadily reduced globally can only be known by also observing or otherwise accounting for the MAC curve of the firm the polluting activity gets shifted to. This analytical leakage problem pertains to MAC curves constructed for any single analytical unit in isolation: a single firm, a single country, a single industry, and to the extent that pollutants are substitutable, a single pollutant. Ignoring this problem gives the misleading impression that for any entity considered in isolation (firm, country, industry, pollutant), the cost of eliminating a unit of pollution ‘in house’ is equivalent to the cost of shifting it elsewhere. The most rigorous way to test the true shape of the MAC curve is therefore to apply the backwards-looking approach to an abatement experience that proceeded simultaneously across all countries, all sectors and all firms. This is a tall order from a data point of view.

In this paper I estimate a MAC-like curve using a dataset from an empirical experience that comes close to satisfying these conditions. The international effort to phase out ozone depleting substances (ODSs) began in earnest in 1989 and had succeeded in reducing ODS consumption by over 85 percent globally by December 2012 relative to a 1990 baseline. Much of this reduction occurred through a financial assistance mechanism established under the Montreal Protocol (MP) whereby relatively rich countries subsidized phase-out in relatively poor countries. A financial record documents the 7,060 projects through which the majority of the ODS phase-out has been achieved across 145 low income countries. The record includes abatement cost and quantity information sufficient to construct a MAC curve, as well as detailed information about the phase-out technology used in many projects.

I use this data to explore the shape of an empirical global MAC curve. I also test explanations related to technology and learning for why the curve takes the shape that it does. The shape of the MAC curve is the subject of considerable disagreement in the literature. Fischer and Morgenstern (2006), Newell and Stavins (2003), Kuik et al (2009) and Tol (2009) depict the curve as being linear or very close. Nordhaus, in his lab notes on the estimation of the curve he used to parameterise the DICE model, describes the most plausible curve as convex 'with a well-determined exponent near 3' (Nordhaus 2007). Fischer and Sterner (2012) find that when the MAC curve is convex, greater uncertainty about the benefit of abatement action implies the need for early abatement. If the curve is concave, uncertainty 'shifts the focus somewhat away from early action'. Practitioners and consultants have suggested that the unit cost of abatement for GHGs over an initial quantity achieved is negative (Enkvist et al 2007). Kesicki and Strachan (2011) and Kesicki and Ekins (2013) refute this, giving several plausible reasons why negative abatement costs are not compatible with an efficient market.[2] This literature generally describes two broad forces that give the MAC curve its shape: (1) the increasing 'difficulty' of reducing each subsequent unit of pollution and the upward pressure this puts on the cost of abatement at the margin and (2) learning and technological advance, which exert downward pressure on the average (Stern 2006) if not also the marginal unit cost.

These two forces feature in prior empirical work on the global phase-out of ODSs. Godwin et al (2010) found that US firms bore substantial cost in replacing ODSs with ozone-benign substances during the US phase-out. This cost was offset to an extent by the up-scaling of production facilities when the firms implemented the phase-out, and by the widening of competition in the production of ODS substitutes across the chemicals industry. DeCanio and Norman (2005) analysed the same project-level data that I analyse here, except their data only ran until to 2003. DeCanio and Norman regressed the cost per ton of ODS eliminated across 1,975 phase-out projects on time, project size, and industry sector dummies. They found that the average project unit abatement cost declined over time. They estimated that the average unit abatement cost declined by US$570 per year between 1991 and 2003. Projects implemented later in time had a lower per-unit abatement cost than projects implemented earlier in time, all else equal. They attribute this to learning by doing and to 'ongoing knowledge gains'.

I find the same initial cost decline that DeCanio and Norman did, but also that cost then increased in subsequent years, forming a convex relationship overall for the total observable phase-out. I describe the policy context and phase-out activities that underlie this pattern in section 2. I describe the technology and learning forces that I suggest gave rise to the convexity in section 3. Results from the empirical model are given in section 4 and section 5 discusses their implications.

---

[2] They argue that what are perceived to be negative abatement costs are explained by: (1) a failure to account for upstream, non-financial and ancillary costs like early stage R&D and transaction costs surrounding implementation of the investment; (2) implementation barriers like the principle-agent problem; and (3) and the difference between the 3.5 percent or lower social discount rate commonly used to estimate a MAC curve and the 20 percent or higher discount rate typically applied by private firms.

2. The global phase-out of ozone depleting substances

The Montreal Protocol on Substances that Deplete the Ozone Layer entered into force on January 1st, 1989. It implemented what was virtually a global ban on ODSs by setting out, for all 197 parties to the agreement, incremental phase-out schedules for nine classes of chemicals. Participation by low-income countries was critical for making the Protocol work (Anderson XXXX). Countries were granted an additional 10 years to achieve phase-out relative to rich countries if (a) they qualified as developing countries and (b) their annual ODS consumption was less than 0.3kg per capita per year. Of the 197 parties to the Protocol, 148 met these criteria and have since adhered to a delayed phase-out schedule. These countries are referred to as ‘Article 5’ countries. Figure 1 compares the phase-out schedules for non-Article 5 and Article 5 countries.

Figure 1: Phase-out schedules agreed to by non-Article 5 and Article 5 countries

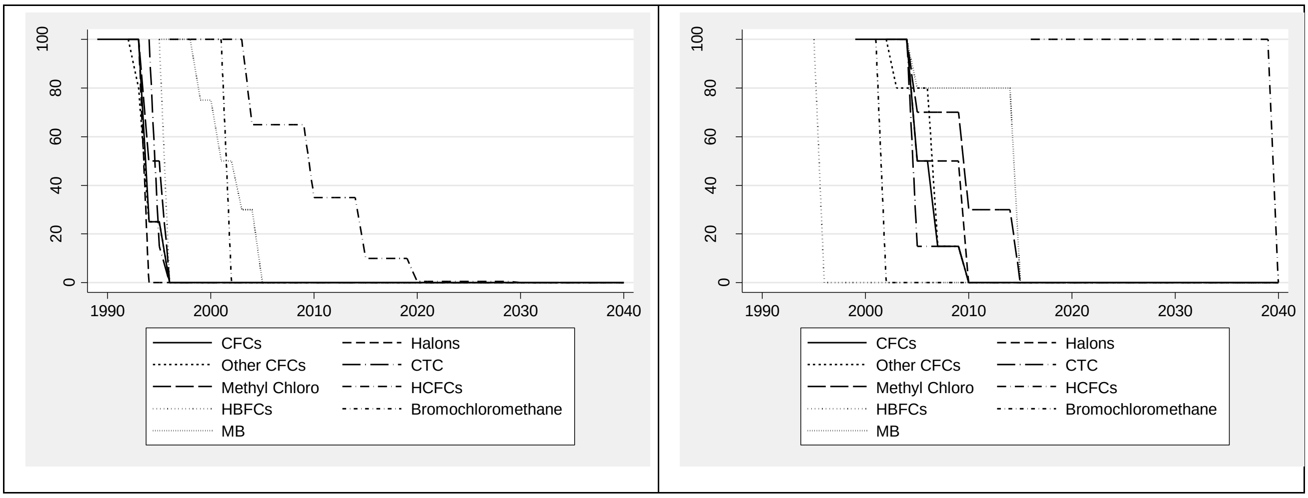

Note: Parties to the Montreal Protocol that qualified for Article 5 status were allowed to implement the ODS ban by a delayed phase-out schedule, relative to non-Article 5 countries. Article 5 countries (right) were obliged to implement the ban around 10 years later for most substances than non-Article 5 countries (left).

The delayed phase-out schedule for Article 5 countries creates quasi natural experiment conditions for identifying a MAC-like curve. The lagged phase-out schedule makes it more certain that the abatement cost being observed in Article 5 countries is the cost of eliminating the pollutant permanently rather than the cost of shifting it elsewhere. If polluting activities had been able to migrate from controlled countries to uncontrolled countries then one might worry that the ‘abatement cost’ observed in the controlled country reflected leakage rather than permanent abatement. This is unlikely to have happened in the ODS phase-out because by the time an Article 5 country began to implement the phase-out, the only countries it could have exported its polluting activities to were either other Article 5 countries that had agreed to the same schedule (in which case the cost of permanent abatement should be observable in the receiving country) or to non-Article 5 (rich) countries which would have been very unlikely to accept those activities since they implemented a ban on them some years earlier.

The Protocol created a financial mechanism to assist Article 5 countries with the cost of phase-out. Article 10 created the Multilateral Fund for the Implementation of the Montreal Protocol (the Multilateral Fund or the Fund) whose purpose was to implement the commitment to meet 'all agreed incremental costs' of phase-out by the Article 5 countries.[3] The Fund began operation in 1991 and has remained in operation continuously to the time of writing. It has supported factory conversion, technical assistance, demonstration, training, planning, and institutional capacity building projects to facilitate phase-out. It is managed by an Executive Committee of 14 members representing equally the interests of donor and recipient countries. Donor countries have replenished the Fund eight times since 1990 with around US$400 million each time. By December 2012 the Executive Committee had met 68 times and approved approximately US$2.6 billion (2012 dollars) in support for projects.

A backwards-looking empirical MAC curve will be more accurate the greater the proportion of total abatement that has been achieved. Only a segment of the curve can be observed if only 10 percent of all the abatement that will ever occur has occurred. Substantial progress has been made toward achieving the global ODS consumption target of nearly zero established under the Protocol. Figure 2 gives the level of ODS consumption between 1990 and 2012 for the ten highest-consuming Article 5 countries. ODS consumption here and throughout is measured in ODP tons.[4] China, whose ODS consumption was four times greater than the next largest country, is omitted from the left panel. By the end of 2012 all ten countries had reduced their ODS consumption by approximately 90 percent relative to 1990 levels.

---

[3] Article 10 reads: 'The Parties shall establish a mechanism for the purpose of providing financial and technical co-operation, including the transfer of technologies, to Parties operating under paragraph 1 of Article 5 of this Protocol to enable their compliance with the control measures . . . . The mechanism, contributions to which shall be additional to other financial transfers to Parties operating under that paragraph, shall meet all agreed incremental costs of such Parties in order to enable their compliance with the control measures of the Protocol.' (UNEP 2000: 33 [Emphasis added])

[4] Different ODSs have different ozone depleting potentials (ODP) in the same way different GHGs have different global warming potentials. The ODP of the reference gas CFC-11 is one but the ODP of halon-2402 is six and the ODP of methyl bromide is 0.6.

Figure 2: ODS consumption for selected Article 5 countries, 1990-2012

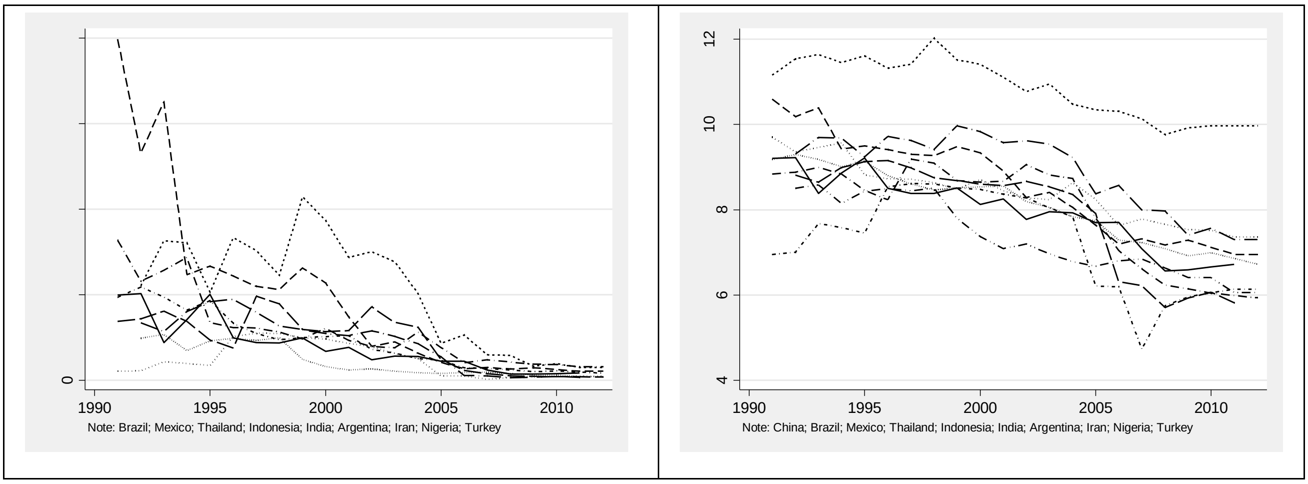

Note: ODS consumption over time by the largest ODS-consuming countries qualifying for Article 5 status. The left panel omits China, the largest consumer.

The phase-out projects supported by the Fund account for a large proportion of this reduction. They do not account for the entire reduction since some phase-out activities were undertaken by Article 5 countries without Fund support. Table 1 shows how much of the total phase-out by Article 5 countries is attributable to Multilateral Fund projects. The first three columns come from ODS consumption level data reported by the 197 parties to the United Nations Environment Program under the Protocol in each year. The last two columns give the aggregate amount of ODS consumption eliminated through the projects. The projects have accounted for 86.5 percent of the reduction achieved by Article 5 countries to date. They have accounted for zero percent of the reduction achieved by non-Article 5 (rich) countries since these countries were not eligible for Fund assistance.

Table 1: Change in country ODS consumption attributable to MF-funded projects

	ODS consumption 1990	**ODS consumption 2012**	**Change 1990-2012**	**Of which MF-assisted**	**Of which MF-assisted %**
Article 5 countries	360,421	40,805	-319,616	-276,615	86.54
Non-Article 5 countries	739,957	3,422	-736,535	0	0
All signatories (total)	1,100,378	44,227	-1,056,151	-276,615	26.19

This means that a large proportion of the Article 5 country phase-out experience is observable in the project data. Assuming conformance with the Protocol the MAC curve estimated below reflects around 86.5 percent of the total ODS abatement that will ever occur across these 148 countries. This does not mean that the curve can be generalised to the entire Article 5 country phase-out experience. This does not mean that it can be generalised to the experience of non-Article 5 countries.

The curve estimated below is a MAC-'like' curve because it is estimated from the average unit cost of abating all the pollution within each project, which I can observe, and not the cost of abating each individual ODS ton at the margin, which I cannot observe.

Phase-out projects for the basis of the empirical analysis. It is important to explain how the Fund chooses which projects to support and how phase-out technologies are chosen. This description is based on extended telephone conversations with two functionaries at the Multilateral Fund who have been closely connected with Fund procedures for selecting projects.

A project is typically proposed by the enterprise that eventually undergoes the phase-out, in collaboration with the country authority responsible for country compliance under the Protocol. The Fund approves and funds project but it does not propose them. Until recently, enterprises have come to the phase-out process more or less voluntarily. Project proposals are often revised several times before they arrive with the Executive Committee for formal approval. Project proposals are frequently completely reformulated in consultation with the Fund Secretariat in weeks 14 to 4 before an Executive Committee meeting. Since the un-fundable features of projects tend to be revised-away through pre-approval negotiation, the Executive Committee receives very few projects that are not approvable. The Fund does not keep records of non-approved projects.

The project described below was approved in 1994 and targeted the refrigeration sector in Algeria for the elimination of 425 tons of ODS-consuming industrial capacity. The project received US$6.5 million of Fund assistance. The primary technology involved was the substitution of CFC-11 for cyclopentane.

> Project description: Replacement of existing low pressure foam machines with high pressure machines, retrofitting of existing high pressure machines, modification of fixtures, acquisition of cyclopentane storage and delivery system, replacement of premixing stations, safety programme (foam component), replacement of the current refrigerant evacuation and charging boards, leak detectors, retrofitting existing vacuum pumps, (refrigerant component), model prototyping, trials and reliability tests. Refrigeration services will be equipped to handle the recovery of both CFC-12 and HFC-134a.

The Fund's remit is to meet the 'incremental cost' of compliance. The project proposer pays for all parts of the project that are not directly connected to ODS elimination. The Fund Secretariat applies a framework of allowable costs to the requirement to reach a technical conversion standard. It funds 100 percent of costs within the framework and funds nothing outside the framework.

Incremental capital and operating costs are allowable. Non-allowable costs include changes to building structures, new equipment unrelated to conversion, new electricity supplies, and R&D. The Fund generally does not fund upgrades to a facility outside the ODS phase-out that would bring it up to the level of international competition. If the Fund anticipates that there will be cost savings from the phase-out changes relative to the pre-phase-out production standard, as a result of the new equipment, the savings are deducted from the approvable project costs. The Fund's contribution to a project is based on the estimated project cost at the time the funding decision is taken. It does not build in expected cost savings that might arise from technological progress, or expected cost increases connected to eliminating later units of polluting capacity. The Fund avoids funding the purchase of equipment that is of a higher quality than necessary to meet the phase-out objective. One of the functionaries mentioned that the standard of approvable equipment tended to be in line 'not with a Mercedes, not with a Volkswagon – maybe with a Skoda'.

Countries must participate in a project-level verification process to receive Fund assistance. All approved projects are implemented under the supervision of an international agency: IBRD, UNDP, UNEP, UNIDO or a country development agency, e.g. USAID. Site visits are conducted by the implementing agency during and after project completion to ensure that old equipment has been destroyed under implementing agency supervision. The purpose of site visits is to prevent old equipment from being resold or reused and to check for illicit back conversion. The risk of back conversion is high for projects where the new operating cost is high and low for projects where the new operating cost is low.

Project proposers are left to choose the phase-out technology insofar as the technology appears on a list of approved technologies maintained by the Fund. The Fund pushes for project proposals to incorporate what it considers to be the 'best possible' phase-out technology. The best possible technology has changed over time. There was major uncertainty about which technologies would be used in the phase-out when the Protocol was signed in 1989. Many of the technologies that eventually emerged involved replacing an ozone depleting substance with an ozone-benign substance or process. The replacement of CFC-11 by HCFC-141b was among the first substitution possibilities to emerge. The Fund avoids funding projects based on technologies that are ozone benign but which might need to be replaced again someday, for example HCFC projects where the global warming potential of the new substance is high.

The Fund Secretariat, the Executive Committee and donor countries to the Fund have all shifted their position over time on the effectiveness of learning-related 'soft' projects. In the early years of the Fund's operation donor countries were largely unwilling to fund country-wide phase-out plans and institutional strengthening measures, because the link to actual phased-out ODS tons was weak or difficult to detect. The Fund Secretariat and Executive Committee came to change their view on this first, followed eventually by the donor countries. It turned out that institutional strengthening projects could appoint competent persons to positions of power who could in turn channel information from site visits to the Fund administrators. Skilled individuals could also support the process of formulating and implementing legislation in the recipient country. This proved crucial for preventing banned equipment and chemicals from 'moving on' to other enterprises that were out of the authorities' direct view. Institutional strengthening projects came to be seen as a way to prevent slide back, to exercise cost control, and to help countries efficiently meet their duties under the Protocol.

3. Data: technology 'possibilities' and 'applications'

Article 5 countries achieved a large amount of their total phase-out through projects funded by the Fund but they also achieved some of the phase-out independently. To model the whole phase-out process and not just the part captured by the projects, I matched two additional datasets to the project data. All countries party to the Protocol were obliged to report their ODS consumption level each year to the UNEP Ozone Secretariat.[5] This is the first additional dataset: the overall ODS consumption level of the country in the year where the project occurred. The second additional dataset includes macroeconomic and development indicator variables from the World Bank, measured for each country in each year. This makes it possible to regress average unit abatement cost observed at the project level on the total amount of ODS consumption remaining in the country at the time the project commenced, and separately, on the underlying technology and learning variables that I suggest give the curve its shape.

I examine all 7,060 phase-out projects supported by the Fund from its inception in 1991 to the end of December 2012. Table 2 below breaks down the projects by type according to the Multilateral Fund's internal project classification system. I further differentiate between two broad project groups based on the information that is available for each project. This is necessary because the cost and quantity information necessary to estimate a MAC curve is only available for some types of projects, and this bears on the modelling approach in the next section.

'Hard' projects involve physically converting factories and production equipment away from ODS-using production methods. This type of project results in directly measurable reductions in ODS consumption. Quantity abated information is provided for these hard (investment) projects.

---

[5] A country's ODS 'consumption' is measured as ODS production, plus imports, minus exports.

‘Soft’ projects by contrast aim to build the capacity of countries and enterprises to achieve phase-out, through country-wide phase-out planning, demonstration activities, institutional strengthening, project preparation, technical assistance, and training. There is no quantity abated information in the data for soft projects. To date, soft projects have outnumbered hard projects but hard projects account for the majority of Fund expenditure.

Table 2: Projects by type

Project type	**Hard / soft**	**Freq.**	**Freq. %**	**Spending (mil. 2012 US$)**	**Spending %**
Country plan	S	168	2.38	10,321	0.32
Demonstration	S	114	1.61	66,468	2.08
Institutional strengthening	S	874	12.38	113,456	3.55
Investment	H	2,792	39.54	2,611,866	81.61
Project preparation	S	1,577	22.33	95,690	2.99
Technical assistance	S	1,202	17.02	267,768	8.37
Training	S	334	4.73	34,859	1.09
Total		**7,060**	**100.00**	**3,200,429**	**100.00**

Note: Fund-funded phase-out projects to December 2012. Quantity abated information is available only for investment projects.

Together, projects funded by the Fund have eliminated approximately 276,600 tons of ODS-consuming industrial capacity to date. Dividing this into the US$3.2 billion in total Fund expenditure gives a global average unit phase-out cost of US$11,568 per ton. In addition to ODS-consuming industrial capacity, projects also phased out large quantities of ODS-producing industrial capacity, mostly at plants in China and India and mostly through plant closure or conversion. In this paper I only consider the cost of phasing out ODS consumption. Conceptualising ‘abatement’ as the sum of consumption and production phased out would introduce double counting and other problems in the estimation of a MAC curve since every ton of ODS produced is also eventually consumed. ODS production projects also occur infrequently (60 out of 7,060) and they are less interesting from a technology point of view because most were dealt with through plant closure.

Information is available about the technology used to achieve phase-out for all hard / investment projects. The phase-out technology is given as the combination of the chemical transitioned away from and the benign substance or substitute transitioned to. In this sense a technology is the technical possibility of substituting one input for another (Orr 1976; Romer 1994). 148 distinct substitution possibilities appear in the entire project record. I refer to each of these as a ‘technology possibility’. Some technology possibilities are applied in more than one project. When this happens I refer to these as separate ‘applications’ of the same technology possibility. The distinction between technology possibilities and applications is discussed below. Table 3 gives the ten most frequently applied technology possibilities and groups the remaining 138 into an ‘other’ category.

Table 3: Most frequently occurring technology possibilities

Technology	**Number of projects applying possibility**	**Percent of projects applying possibility**	**ODS tons phased out by projects using possibility**
CFC-11 to HCFC-141b	791	27.17	31,308
CFC-12 to Investment/non-investment activities	203	6.97	5,484
CFC-11 to Water/carbon dioxide	170	5.84	5,769
CFC-11 to Methylene chloride	158	5.43	8,022
CFC-12 to Multiple actions	140	4.81	10,230
CFC-12 to Recovery/recycling	138	4.74	6,487
CFC-11 to Cyclopentane	132	4.53	19,420
Methyl bromide to various alternative technologies	71	2.44	4,495
CFC-12 to Hydrocarbon aerosol propellant	65	2.23	18,149
CFC-12 to Recovery/recycling/retrofit	60	2.06	1,872
Other	983	33.55	165,381
Total	**2,911**	**100.00**	**276,615**

Note: a technology here is the possibility to replace an ozone-depleting substance with a benign substitute. When a technology possibility is used in more than one project each use is an application of the possibility. One might expect the unit cost of abatement to decline with repeat applications within a technology possibility.

All of the 148 technology possibilities involved transitioning away from one of the controlled substances under the Protocol and most involved transitioning to a benign chemical substance. However some involved transitioning to something other than a benign chemical substance, for example a practice, process, method, regime, behaviour or procedure that achieved the same result as the controlled substance. The second most frequently applied technology possibility ‘CFC-12 to Investment/non-investment action’ in Table 3 is one of these non-chemical substitutes. Examples of projects that substituted knowledge, behaviour or social practice for a controlled substance are: a project in Peru in 2004 that sought to prevent illegal CFC trade by providing technical assistance to border agency officials; a project in Benin in 2000 that trained officials to investigate and inspect refrigeration equipment; and a project in China in 2009 that sought to eliminate fugitive ODSs by establishing a network of domestic appliance dismantling stations for refrigerant recovery.

These 'institutional' or 'organisational' technology possibilities were the primary technology for around 7 percent of all hard / investment projects. Discussions of technology in large complex pollution challenges like mitigating GHG emissions usually only relate to 'technology' in the sense of physical machines or devices. MAC curves that depict behaviour change as a reduction method are rare. Table 3 shows that these technologies played a significant role in the global ODS phase-out by low-income countries. In policy discussions about the role of 'environmental technology' in dealing with future pollution challenges like GHG mitigation, it would therefore be realistic to include abatement possibilities arising from behavioural, procedural and occupational changes.

As above I distinguish between a technology possibility and an application of a technology possibility and expect these to exert different influences on the shape of the MAC curve. There is a strong empirical basis for believing that new technology possibilities emerge when the cost of whatever input the technology saves, rises (Brunnermeier and Cohen 2003; Hicks 1932; Popp 2002). There is also a strong basis for believing that repeat applications within a given technology possibility bring down the cost of installing the possibility at the project level and/or bring down the unit cost of abatement achieved by the possibility (Dutton and Thomas 1984; Gillingham and Bollinger 2012; McDonald and Schrattenholzer 2001). I have no compelling reason to expect the emergence of a new technology possibility to associate with lower abatement cost since theory suggests that the cost-lowering mechanism is repeat applications within technology possibilities.

The emergence of a new technology possibility is still an important event, even if this event is not in itself cost-lowering. To elaborate the relationship with abatement cost I ordered the 148 technology possibilities in the sequence in which they emerged, on the basis of the month in which each technology possibility was applied for the first time. The left panel of Figure 3 plots the rank of the emergence of each possibility against mean unit cost of all the projects using the possibility. The left panel of Figure 3 shows that early-emerging possibilities associate with lower unit cost.

Figure 3: Abatement cost and project size v rank of emergence of technology possibility

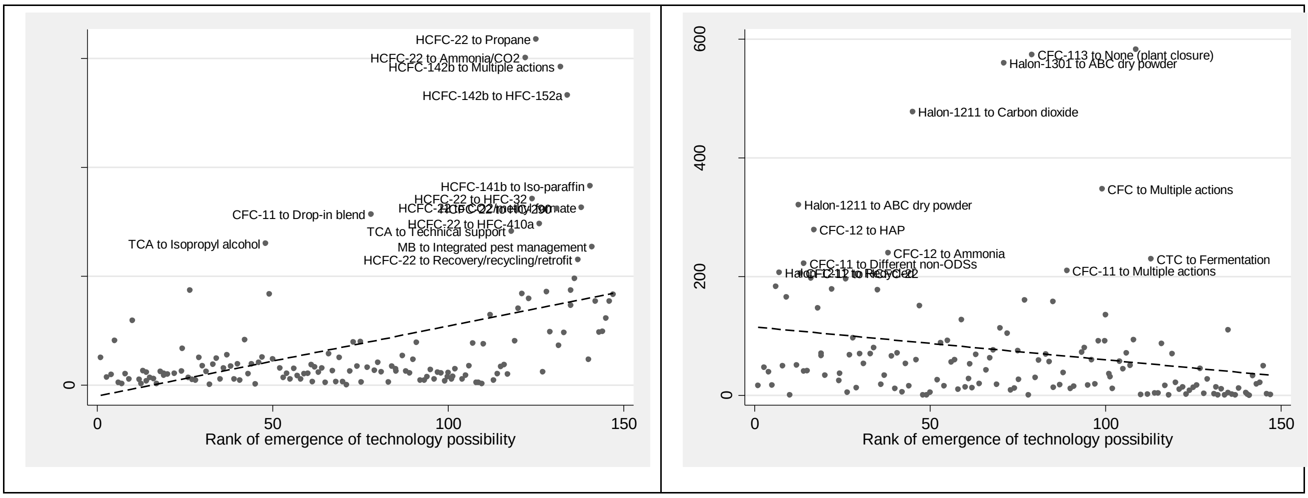

Note: early-emerging technology possibilities associate with lower unit abatement cost (left) and with a larger average project size as measured by tons of ODSs phased-out. Points are technology possibilities ordered by the month of first application. N=148 in both.

This is consistent with the idea that early-emerging possibilities were transplanted into the realm of the ODS abatement effort from other sectors outside ODS abatement where they were already technically proven and in use. Transplanting, e.g. re-purposing possibilities from pre-existing uses would have involved little or no discovery cost since the blueprint was already in existence and the possibility just needed to be adapted and possibly scaled-up. The flip side of this explanation is that late-emerging possibilities would have needed to be invented. These would have involved dealing with specific chemicals or industrial processes for which no easily adaptable, pre-existing possibility existed. They therefore would have involved higher knowledge search, discovery and application costs which were eventually reflected in higher unit abatement cost.

The right panel in Figure 3 plots the rank of the emergence of each technology possibility against the quantity abated using possibility. Quantity abated is measured as the average size of all the projects that used the possibility. Whereas early-emerging possibilities associate with lower unit abatement cost in the left panel, they associate with larger quantities of pollution abated in the right panel. This implies that the first possibilities to emerge were the ones with the biggest immediate deployment potential. They may have had low development costs because they were being transplanted or adapted from pre-existing blueprints. They may also have had large potential abatement markets since the chemicals and industrial processes they could achieve abatement for were widely used. Later-emerging possibilities by contrast might be searched-out, adapted or invented at relatively high cost to meet abatement needs in small, niche, low-volume areas where no other possibility was suitable.

Figure 4 conveys how the emergence of new technology possibilities on the one hand, and repeat applications within each possibility on the other, can play different roles in giving shape to an overall MAC curve. TP_1 denotes the first technology possibility to emerge while TP_1,A_1 is the first application of the possibility in the form of an abatement project. Cost-lowering, within-technology learning causes the marginal cost to fall until abatement possibilities suitable to TP_1 have been exhausted. Technology possibilities that emerge subsequently tend to be more expensive on average, but with learning effects driving down the marginal cost over repeat applications within each. The relationship between each of these forces and average project unit abatement cost is estimated in the next section.

Figure 4: Technology possibilities and repeat applications in a MAC curve

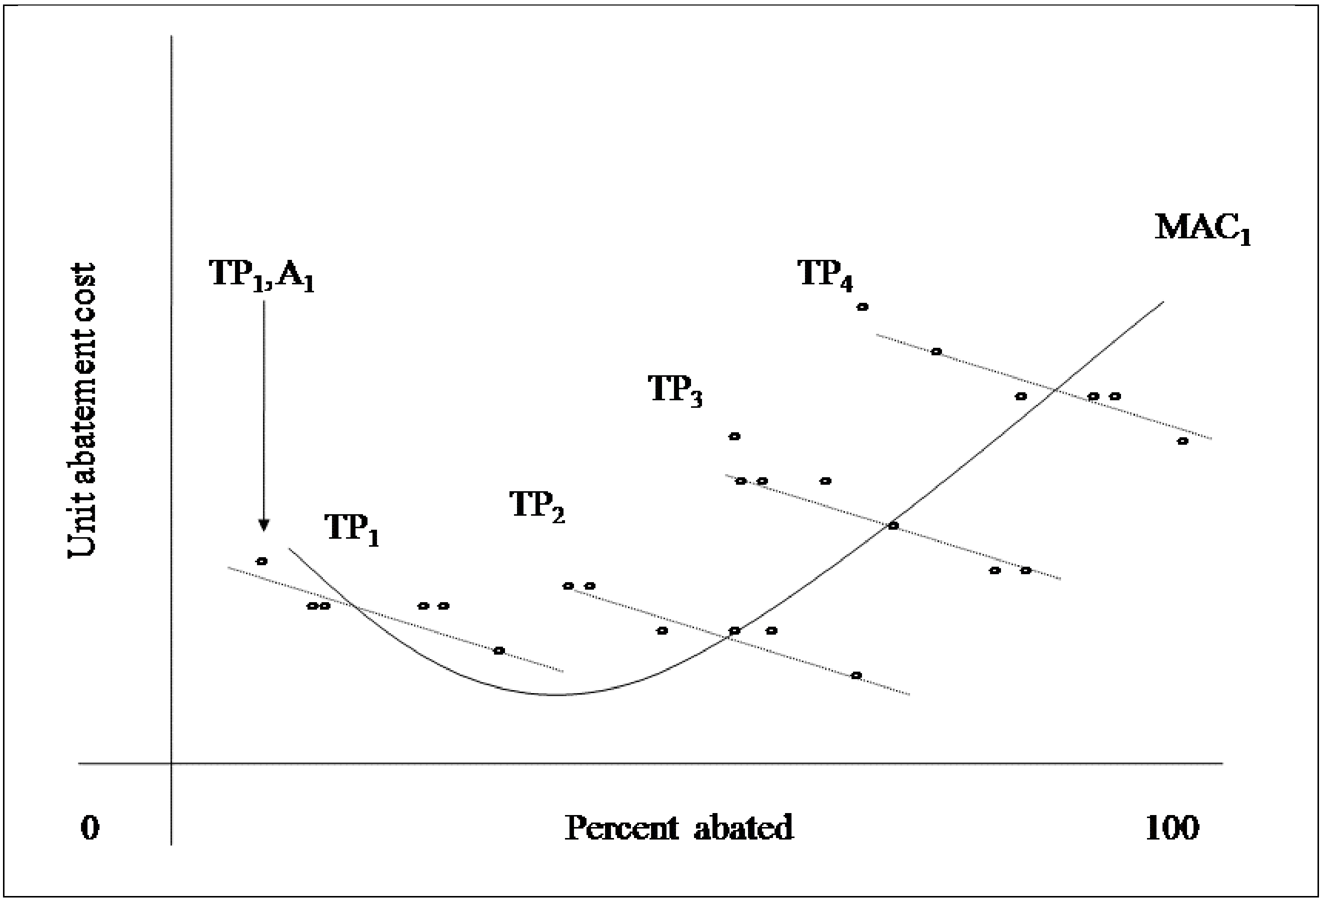

Note: A MAC curved comprised of two forces: the emergence of new technology possibilities (TP) to deal with increasingly complicated and expensive abatement problems, and repeat applications within those possibilities (A) in the form of projects for example that drive down marginal cost through learning.

4. Regressions

The regressions estimate the relationship between the average unit cost of abatement at the project level and the quantity of abatement that had been achieved in the country in the year the project was initiated. Various functional forms are tested to estimate the best-fitting shape of this relationship. Technology possibility and repeat application variables are then tested in place of the quantity abated variable to see if these underlying learning and technology forces can explain the abatement cost change pattern. This attempts to identify a deeper mechanism for the changing cost of abatement.

a. Variables

The dependent variable in all specifications is the average project unit abatement cost, which is measured at project level, i. The numerator is total project cost and the denominator is tons of ODSs phased out by the project. Superscript hard indicates that the observations that comprise the dependent variable are hard (investment) projects only since these are the projects for which quantity abated information is available.

$$Av. project\ unit\ abatement\ cost\,_i^{hard} = \frac{Cost\,_i^{hard}}{Tons\ ODS\ abated\,_i^{hard}}$$

The independent variable of initial interest is the percent ODS abated by each country in each year. To calculate this, I transformed the original ODS consumption level values, which were measured in ODP tons for each country-year and described in Figure 2, so that 100 corresponds to a country's maximum-ever ODS consumption level and zero to zero consumption.

$$ODS\ cons. level\,_{c,t}^{normalised} = 100 * \left(\frac{ODS\ cons. level\,_{c,t}}{ODS\ cons. level\,_c^{max}}\right)$$

This re-scaling of the original values gives a within-country maximum pollution level that is comparable across countries. To move from pollution level to an 'abatement achieved' interpretation I then reversed the 0 – 100 scale. Zero now corresponds to a country-year where the country is at its maximum-ever pollution level and no abatement has occurred, and 100 to a country-year where pollution is at its minimum and all observable abatement has occurred.

$$ODS\ abatement\ achieved\,_{c,t}^{normalised} = 100\ -\ ODS\ cons.\ level\,_{c,t}^{normalised}$$

The variable now measures the extent of abatement activity in a way that is comparable across large and small countries. The amount of abatement achieved by China and Palau needs to be comparable because abatement cost is expected to rise the closer quantity abated gets to 100 percent within each country and because countries were the entities responsible for compliance under the Protocol.

Technology possibility sequence is the order in which each technology possibility appeared for the first time in a Fund-funded project. This can be thought of as the rank order of appearance of each of the 148 technology possibilities. For example, the first possibility to emerge was ‘CFC-113 to Semi-aqueous cleaning’. It appeared for the first time in a relatively small, solvent sector project in Mexico in 1991. The 148th technology possibility to emerge was ‘HCFC-22 to Non-HCFC’. It appeared for the first time in December 2012 in a technical assistance project in Thailand. This variable tests the expectation that later-appearing technology possibilities emerge under higher abatement cost conditions.

Possibility application sequence is the order in which each application of each technology possibility appeared. This is measured as the rank order by time of the projects that applied the possibility, within each possibility. This means that the first application of the technology possibility ‘CFC-113 to Semi-aqueous cleaning’ was in the same project where the possibility itself appeared for the first time: in Mexio in 1991. The last application of the possibility (it was applied in 10 projects) happened in a solvent sector project in China in 1997. Capturing the order of repeat applications within each technology possibility tests whether the unit cost of abatement declines with repeat applications. The expectation is that repeat applications of a technology should be cost-lowering as agents accumulate experience with it (Arrow 1962; Mann and Richels 2004).

Control variables are as follows.

The scale of the project is expected to influence unit cost. The variable scale measures the number of months the project was on-going, a proxy for size.

Soft project spending should create abatement-relevant technical knowledge that influences the unit abatement cost observed in hard / investment projects. The stock of soft technical knowledge is measured as cumulative soft project spending in each country over the 22 years, depreciated by 10 percent annually.

147 country dummies account for all observable and unobservable, time-invariant country-specific factors. 11 industry sector dummies[6] do the same, the categories having been defined by the Multilateral Fund. 21 year dummies control for time-varying factors.

Country, sector, and year fixed effects mean that if omitted variables are biasing the coefficients of interest then time-varying influences must be at play. I use four time-varying variables as checks: GDP per capita (constant 2005 USD), GDP growth rate (percent), industry value added (percent of GDP), and trade (percent of GDP).

Table 4 gives descriptive statistics for all continuous variables.

---

[6] Aerosols, desiccants, foams, fumigants, halons, multisector, other, process agent, refrigerant, solvent, sterilant.

Table 4: Descriptive statistics

	Units	Varies by	N	Mean	Min	Max	SD
Av. project unit abatement cost	Thou. 2012 USD per ton	i	2,911	28.00	0.00	4,718.60	113.07
ODS abatement achieved	Proportion	c,t	1,863	56.51	0.00	100.00	32.82
Tech. possibility sequence	Rank	i	2,911	40.51	1.00	148.00	40.97
Possibility application sequence	Rank	i	2,911	140.81	1.00	792.00	200.72
Scale	Months	i	2,911	37.93	0.00	195.00	121.12
Soft project spending	Thou. 2012 USD	c, t	2,911	7,432.22	67.11	75,852.07	8,842.51
Country income	GDP per cap.	c,t	1,834	3,552.83	118.64	121,340.23	6,533.20
Country growth	GDP growth	c,t	1,843	4.25	-32.83	34.39	4.42
Country industry VA	Perc. of GDP	c,t	1,793	29.52	4.22	78.2	12.31
Country trade	Perc. of GDP	c,t	1,810	78.65	0.36	223.06	38.06

b. Estimations

I initially regress average project unit abatement cost on ODS abatement achieved, scale, a vector of control variables (soft project spending and the country economic performance variables), and a vector of 147 country dummies.

$$\begin{aligned} &Av.project\ unit\ abatement\ cost\,_i^{hard} \\ &\qquad = a_i + \text{â}_1(ODS\ abatement\ achieved)_{c,t}^{normalised} + \text{â}_2(Scale)_i + {}'\text{ë}_{c,t} + {}'\text{ä}_c \\ &\qquad + \text{å}_i \end{aligned}$$

where delta is a vector of country dummies and ε a classical error term. The aim is to estimate the coefficient β_1, the slope between average unit abatement cost and abatement achieved. The data are structured as a pooled cross section with country fixed effects. Various attempts were made to structure the data as a panel, which would have made it possible to test the effect of the emergence of technology possibilities on abatement cost in subsequent periods, but there were too many gaps in the sequence of years for many countries. Reverting to a pooled cross section of projects, the model is estimated by OLS with robust standard errors.

In column (1) of Table 5, more ODS abatement achieved by a country associates with higher average unit abatement cost at the project level as expected. The coefficient β_1 implies that a one percent increase in quantity abated by a country from its maximum-ever peak associates with an increase in the average unit abatement cost of US$399 in real terms. Columns (2) – (5) attempt to undermine this relationship first with industry sector dummies, then with year dummies, then with country and year dummies. Statistical significance falls away when country and year fixed effects are included together in specification (4). When the time-varying country development variables are added in specification (5) the coefficient of interest is positive, significant and of roughly the same magnitude as in specification (1). No omitted variable bias has been flagged up by the time-varying variables and I have no good theoretical reason to include them in further specifications. The results in Table 5 are interpreted as weak evidence of a linear MAC-like relationship.

Table 5: MAC curve: quantity abated explanation

	(1)	(2)	(3)	(4)	(5)
Country abatement achieved	0.391***	0.187***	0.076*	0.099	0.393***
	(0.064)	(0.043)	(0.037)	(0.070)	(0.066)
Scale, project length	0.094	0.207*	0.287**	0.188	0.090
	(0.106)	(0.084)	(0.100)	(0.098)	(0.103)
Soft project spending	0.001**	0.001	0.000	0.001	0.001**
	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)
Country income					0.001
					(0.005)
Country growth					0.460
					(0.329)
Country industry VA					0.382
					(0.723)
Country trade					-0.194
					(0.170)
Constant	-13.763*	-10.230	4.229	16.304	-8.544
	(5.645)	(5.502)	(9.523)	(13.335)	(21.953)
Country dummies	Yes	No	No	Yes	Yes
Industry dummies	No	Yes	No	No	No
Year dummies	No	No	Yes	Yes	No
Obs	2911	2911	2911	2911	2850
R-square	0.103	0.0356	0.0833	0.155	0.100
Adj. R-square	0.0572	0.0312	0.0757	0.105	0.0556
Df	135	12	24	157	129

Dependent variable: Average project unit abatement cost, thou. USD per ton.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 6 below explores the functional form of the relationship. Specification (1) is the lin-lin baseline specification. Robustness checks show that the residuals are mildly right skewed, the significance of the relationship of interest is sensitive to outliers (very expensive projects), and the predicted values are poor predictors of observed abatement cost. Specification (2) in Table 6 is a log-lin specification which has similar issues to (1), in particular weak predictive power. Specification (3) adds year dummies to the log-lin specification and again the significance of the relationship falls away.

Specification (4) in Table 6 adds the second order polynomial transformation of the country abatement achieved variable to test whether the relationship of interest might be U-shaped. Evidence of a U-shaped relationship would be consistent DeCanio and Norman’s (2005) finding, which was based on the same data used here until 2003 only, that average project abatement cost declined with the passage of time. It would also be consistent with Kesicki and Strachan’s (2011) and Kesicki and Ekins’s (2013) reasoning that abatement at negative cost is incompatible with an efficient market. Specification (4) gives some evidence of a U-shaped relationship but this does not survive the addition of year dummies in specification (5).

Table 6: Functional form tests

	(1)	(2)	(3)	(4)	(5)
Country abatement achieved	0.391^{***}	0.004^{***}	0.001	-0.516^{***}	0.314^{*}
	(0.064)	(0.001)	(0.001)	(0.149)	(0.160)
Scale, project length	0.094	-0.001	0.000	0.124	0.189
	(0.106)	(0.001)	(0.001)	(0.103)	(0.099)
Soft project spending, stock	0.001^{**}	0.000^{**}	0.000	0.001^{**}	0.001
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Country abatement achieved (^2)				0.010^{***}	-0.003
				(0.002)	(0.003)
Constant	-13.780^{*}	1.467^{***}	2.221^{***}	2.928	10.102
	(5.648)	(0.063)	(0.331)	(5.893)	(16.497)
Country dummies	Yes	Yes	Yes	Yes	Yes
Year dummies	No	No	Yes	No	Yes
Obs	2911	2911	2911	2911	2911
R-square	0.103	0.189	0.347	0.107	0.156
Adj. R-square	0.0572	0.147	0.308	0.0611	0.105
Df	136	136	157	137	157

Dependent variable: Average project unit abatement cost, thou. USD per ton. DV is logged in specifications (2) and (3). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

In the next set of regressions I move beyond the amount of pollution yet to be abated as the main explanation for the observed abatement cost distribution, to a more causal hypothesis that relies on technology and learning.

The hypothesis centres on the idea that new technology possibilities were ‘recruited’ from the realm of technical knowledge that already existed in the global chemicals and allied industries, to the realm of knowledge for abating ODSs. In this explanation, the Multilateral Fund, its donors, and enterprises needing to phase out their ODS consumption began searching for existing technology possibilities that would be capable of bringing down the cost of abatement. Their search was initially fruitful: lots of technology possibilities that could be adapted to the ODS phase-out problem already existed elsewhere and just needed to become known to this group trying to solve the new problem of abatement. Recruiting these technologies into the abatement knowledge realm brought down the cost of abatement. So did repeat project applications within each technology possibility through learning and refinement. Eventually the search for new technology possibilities dried up however and enterprises still obliged to do abatement either had to invent universally new technology possibilities rather than just transplant pre-existing ones (Antonelli 2002; Grover 2013), which was expensive, or abate with inappropriate technology, which was also expensive.

This hypothesis rests on the initial abundance and steady exhaustion of technology possibilities for dealing with pollution rather than on the ‘difficulty’ of the remaining pollution itself. Rather than ‘low hanging’ possibilities drying up, e.g. those that were technically achievable and relatively inexpensive, what dried up was the pre-existing, transferrable, and therefore inexpensive technology possibilities for doing the abatement. The novelty of the hypothesis is the idea that abatement cost changes most basically in accordance with the flow of new technology possibilities into the realm of understanding of those performing the abatement.

I test this hypothesis by dropping the country abatement achieved variable which has been the test variable in all previous specifications. This variable belongs in the model when theory predicts that abatement cost changes because pollution units become more scarce or ‘difficult’ to abate. The hypothesis however attributes changing abatement cost to the availability of pre-existing technology possibilities initially, the recruitment of these from one knowledge realm to another, the exhaustion of pre-existing possibilities, and eventually the hard and expensive work of inventing universally ‘new’ possibilities.

The estimations now includes as test variables technology possibility sequence, its second order polynomial, and the application (project) sequence within each possibility. Delta represents country dummies and theta year dummies.

$$\begin{aligned} Av.project\ unit\ abatement\ cost_{i}^{hard} &= a_i + \hat{a}_1(Tech.possibility\ sequence)_{c,t}^{normalised} \\ &+ \hat{a}_2(Tech.possibility\ sequence^2)_i + \hat{a}_3(Possibility\ application\ sequence)_i \\ &+ \hat{a}_4(Scale)_i + \hat{a}_5(Soft\ project\ spending)_{c,t} + {}'\ddot{a}_c + {}'\grave{e}_t + \mathring{a}_i \end{aligned}$$

In Table 7, specification (1) includes technology possibility sequence on its own with the controls and country fixed effects. Each newly emerging technology possibility associates with a roughly US$400 increase on average in the cost of phasing out a ton of ODS. Specification (2) adds the second order polynomial of this variable. These coefficients imply that while later-emerging possibilities associated with higher abatement cost, early emerging possibilities associated with lower cost, e.g. that the relationship between the appearance of new technology possibilities and abatement cost is convex. Specification (3) adds the application sequence testing whether abatement cost declined within technology possibilities through repeat applications, e.g. projects. Each subsequent application associates with a US$27 decrease in abatement cost on average, holding all else constant. These relationships remain statistically significant when year dummies are included in specification (4).

Table 7: MAC curve: technology and learning explanation

	(1)	(2)	(3)	(4)
Tech. possibility sequence	0.400^{***}	-1.530^{***}	-1.710^{***}	-1.037^{***}
	(0.042)	(0.187)	(0.202)	(0.187)
Scale, project length	0.078	0.218^{*}	0.204^{*}	0.175
	(0.108)	(0.104)	(0.102)	(0.095)
Soft project spending, stock	0.001^{*}	0.001	0.001	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Tech. possibility sequence (^2)		0.017^{***}	0.018^{***}	0.012^{***}
		(0.002)	(0.002)	(0.002)
Possibility application sequence			-0.027^{***}	-0.027^{***}
			(0.006)	(0.007)
Constant	-39.629^{***}	-19.930^{***}	-11.252^{*}	4.425
	(5.087)	(5.006)	(4.980)	(12.213)
Country dummies	Yes	Yes	Yes	Yes
Year dummies	No	No	No	Yes
Obs	2911	2911	2911	2911
R-square	0.112	0.145	0.147	0.165
Adj. R-square	0.0664	0.101	0.102	0.115
Df	136	136	138	159

Dependent variable: Average project unit abatement cost, thou. USD per ton
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

5. Discussion and conclusions

I used project-level data from the Multilateral Fund for the Implementation of the Montreal Protocol to test various explanations for the pattern in the cost of abating ODSs, across 145 low income countries and accounting for around 86.5 percent of the phase-out achieved by these countries between 1990 and 2012. I find that technology and learning forces explain the observable abatement cost pattern better than the rather deterministic ‘quantity remaining to be abated’ explanation and the intrinsic ‘difficulty’ of the later units of pollution that this explanation implies. I considered the argument that initially the cost of dealing with pollution can be negative (Enkvist et al 2007), the weaknesses of that argument (Kesicki and Strachan 2011, Kesicki and Ekins (2013), and previous research on the MAC curve shape. I find that the cost of abating ODSs declined before it started to rise, consistent with what DeCanio and Norman (2005) found using a portion of the same data.

The possibility should not be dismissed that zero or negative cost abatement did actually occur in the global phase-out of ODSs, but that this abatement is not observable in the Multilateral Fund data. Projects to close loss-making factories would not have appeared because they would not have been eligible for 'incremental' support since the project cost was zero. Unobservable zero-cost projects would pull down the real cost curve in the initial period. Fund-funded projects did however account for the large majority of abatement achieved in these countries (86.5 percent) and while the whole picture is not observable, a large part of it is. I have also suggested theoretical reasons for why technology and learning forces might cause the cost of abatement to decline initially before rising, in a convex relationship with technology and learning.

The technology and learning hypothesis shifts the weight of explanation away from some units of pollution being more 'difficult' to abate than others, towards the idea that difficulty is a function of the technology possibilities that are available to the agent performing the abatement. In a way that I have not seen done in the empirical literature before, I distinguish between (1) the appearance of technology 'possibilities' for doing abatement and (2) the repeat 'application' of those possibilities in the form of individual projects. I find evidence to support the idea that the transplantation of pre-existing technology possibilities to the new problem of dealing with ODSs initially lowered the cost of abatement, but when pre-existing possibilities become exhausted and costly knowledge search and invention of new possibilities became necessary, the unit cost of abatement increased. Holding the appearance of new possibilities constant, I find that each subsequent application of each technology possibility associates with a lower unit cost of abatement as would be expected from a learning by doing effect.

These relationships were investigated in a static estimation framework that clearly omits some important dynamic complexities. For example, the activities observed here as 'soft' projects which were mainly focused on human capital building, probably lowered the cost of dealing with pollution with a lag. Attempts to structure the Fund data as a panel to explore these relationships were frustrated by gaps across years within countries.

If the convex relationship between ODS abatement cost and technology and learning forces observed here were real, and if this relationship were generalisable to the problem of reducing GHG emissions, then technology and learning advances could provide one explanation, of many, for why the traded permit price under the EU ETS has been declining and/or so persistently low.

6. References

Antonelli, C (2002) The Economics of Innovation, New Technologies and Structural Change. London: Routledge.

Brunnermeier, SB and MA Cohen (2003) 'Determinants of environmental innovation in US manufacturing industries', Journal of Environmental Economics and Management 45: 278-293.

Carraro, C, E De Cian, L Nicita, E Massetti and E Verdolini (2010) 'Environmental Policy and Technical Change: A Survey', International Review of Environmental and Resource Economics 4: 163–219.

Copeland, BR and MS Taylor (2004) 'Trade, Growth and the Environment', Journal of Economic Literature 42(1): 7-71.

DeCanio, SJ and CS Norman (2005) 'Economics of the "critical use" of methyl bromide under the Montreal Protocol', Contemporary Economic Policy 23 (3): 376-393.

Dutton, JM and A Thomas (1984) 'Treating Progress Functions as a Managerial Opportunity', The Academy of Management Review 9 (2): 235-247.

Enkvist, PA, T Nauclér and J Rosander (2007) 'A cost curve for greenhouse gas reduction: A global study of the size and cost of measures to reduce greenhouse gas emissions yields important insights for businesses and policy makers', McKinsey Quarterly 1: 35-45.

Fischer, C and RD Morgenstern (2006) 'Carbon Abatement Costs: Why the Wide Range of Estimates?', The Energy Journal 2: 73-86.

Fischer, C and T Sterner (2012) 'Climate Policy, Uncertainty and the Role of Technological Innovation', Journal of Public Economic Theory 14(2): 285-309.

Fisher, AC and U Narain (2003) 'Global Warming, Endogenous Risk, and Irreversibility', Environmental and Resource Economics 25: 395–416.

Gillingham, K and B Bollinger (2012) ‘Peer Effects in the Diffusion of Solar Photovoltaic Panels’, Marketing Science 31(6): 900-912.

Gillingham, K, RG Newell and WA Pizer (2008) ‘Modeling endogenous technological change for climate policy analysis’, Energy Economics 30: 2734-2753.

Godwin, DS, MM Van Pelt and TL Krasney (2010) ‘An analysis of reduction opportunities for consumption of hydrofluorocarbons and comparisons to US climate policy proposals’, Journal of Integrative Environmental Sciences 7 (S1): 187-199.

Grover, D (2013) 'The "advancedness" of knowledge in pollution-saving technological change with a qualitative application to SO2 cap and trade', Ecological Economics 89: 123-134.

Grubb, M, C Carraro and J Schellnhuber (2006) ‘Technological Change for Atmospheric Stabilization: Introductory Overview to the Innovation Modeling Comparison Project’, Energy Journal, Special Issue: Endogenous Technological Change and the Economics of Atmospheric Stabilisation, 1-16.

Hicks, JR (1932) The Theory of Wages. London: Macmillan.

Intergovernmental Panel on Climate Change (2013) ‘Agriculture’, in Climate Change 2007: Mitigation. Contribution of Working Group III to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change [B Metz, OR Davidson, PR Bosch, R Dave, LA Meyer (eds)]. Cambridge: CUP.

Kesicki, F and P Ekins (2012) ‘Marginal abatement cost curves: a call for caution’, Climate Policy 12(2): 219-236.

Kesicki, F and N Strachan (2011) ‘Marginal abatement cost (MAC) curves: confronting theory and practice’, Environmental Science & Policy 14(8): 1195-1204.

Klepper, G and S Peterson (2006) ‘Marginal Abatement Cost Curves in General Equilibrium: The Influence of World Energy Prices’, Resource and Energy Economics 28: 1-23.

Kolstad, CD (1996a) ‘Fundamental Irreversibilities in Stock Externalities’, Journal of Public Economics 60: 221–233.

Kolstad, CD (1996b) ‘Learning and Stock Effects in Environmental Regulation: The Case of Greenhouse Gas Emissions’, Journal of Environmental Economics and Management 31: 1–18.

Kuik, O, L Brander and RSJ Tol (2009) ‘Marginal abatement costs of greenhouse gas emissions: A meta-analysis’, Energy Policy 37: 1395-1403.

Magat, W (1978) ‘Pollution control and technological advance: A dynamic model of the firm’, Journal of Environmental Economics and Management 5 (1): 1-125.

McDonald, A and Schrattenholzer, L (2001) ‘Learning rates for energy technologies’, Energy Policy 29: 255-261.

Metcalf, G and D Weisbach (2008) ‘The Design of a Carbon Tax’, NBER Working paper 14375.

Newell, RG and RN Stavins (2003) ‘Cost heterogeneity and the potential savings of market-based policies’, Journal of Regulatory Economics, 23(1): 43–59.

Nordhaus (1991) ‘A Sketch of the Economics of the Greenhouse Effect’, The American Economic Review 81 (2): 146-150.

Nordhaus, W (2007) ‘Accompanying Notes and Documentation on Development of DICE-2007 Model: Notes on DICE-2007.delta.va as of September 21, 2007.’ Yale University, October 5, 2007. www.econ.yale.edu/~nordhaus/homepage/Accom_**Notes**_100507.pdf.

Orr, L (1976) ‘Incentive for Innovation as the Basis for Effluent Charge Strategy’, The American Economic Review 66 (2): 441-447.

Popp, D (2002) ‘Induced innovation and energy prices’. The American Economic Review 92 (1): 160-180.

Popp, D (2010) ‘Exploring Links Between Innovation and Diffusion: Adoption of NOx Control Technologies at US Coal-fired Power Plants’, Environmental and Resource Economics 45: 319-352.

Pye, S, K Fletcher, A Gardiner, T Angelini, J Greenleaf, T Wiley, H Haydock (2008) ‘Review and Update of UK Abatemetn Costs Curves for the Industrial, Domestic and Non-Domestic Sectors’, AEA Energy & Enviroinment, Didcot.

Romer, P (1994) ‘The Origins of Endogenous Growth’, Journal of Economic Perspectives
8 (1): 3–22.

Stern (2006) ‘Stern Review on the Economics of Climate Change (pre-publication edition)
Executive Summary.’ London: HM Treasury.

Tol, RSJ (2009) ‘Intra-union flexibility of non-ETS emission reduction obligations in the European Union. Energy Policy, 37(5): 1745 – 1752.

Vijay, S, JF DeCarolis and RK Srivastava (2010) ‘A bottom-up method to develop pollution abatement cost curves for coal-fired utility boilers’, Energy Policy 38(5): 2255-2261.

Wigley TML, Richels R, Edmonds JA (1996) ‘Economic and environmental choices in the stabilisation of atmospheric CO2 concentrations’, Nature 379:242–245.