Carbon tax or carbon permits: the impact on generators' risks

Richard Green*

Institute for Energy Research and Policy
University of Birmingham
Birmingham B15 2TT
Tel: +44 121 415 8216
Email: r.j.green@bham.ac.uk

March 2007

Abstract

Volatile fuel prices affect both the cost and price of electricity in a liberalised market. Generators with the price-setting technology will face less risk to their profit margins than those with a technology that is not price-setting, even if its costs are not volatile. Emissions permit prices may respond to relative fuel prices, further increasing volatility. This paper simulates the impact of this on generators' profits, comparing an emissions trading scheme and a carbon tax against predictions for the UK in 2020. The carbon tax reduces the volatility faced by nuclear generators, but raises that faced by fossil fuel stations. Optimal portfolios would contain a higher proportion of nuclear plant if a carbon tax was adopted.

JEL: L94
Keywords: Electricity, emissions trading, emissions taxes, fuel price risk

* This research is supported by the Engineering and Physical Sciences Research Council through its Supergen FutureNET consortium, grant GR/S28082/01. I hope that my intellectual debt to the researchers of the Cambridge Electricity Policy Group for many of the concepts used in this paper is clear, and I would like to thank them for hospitality at conferences, and on sabbatical. I would also like to thank participants at the British Institute for Energy Economics / UKERC conference, September 2006, the IDEI Conference on Energy Markets, January 2007, and at seminars at the University of Birmingham and the Competition Commission, for helpful comments, subject to the usual disclaimer.

1. INTRODUCTION

Emissions trading schemes and carbon taxes work by raising the cost of producing electricity from fossil fuels. The more polluting the technology, the more its variable cost increases. This should provide an incentive to shift generation towards low-carbon technologies. We know that taxes and permit schemes should have equivalent results in a world of certainty, but that they perform differently in an uncertain world. This paper considers their different impact on the profit risk faced by electricity generators with a choice of technologies.

The standard analysis of investment choices in electricity uses a cost-minimising approach. Investors do not choose to build new power stations solely on the basis of their expected costs, however. Companies wish to make profits, and to avoid excessive risks. Roques *et al.* (2006b) show that a probabilistic analysis is needed to give the full picture of the expected return on an investment, and its distribution. One particular issue discussed by Roques *et al.* (2006c) is that if the cost of gas and carbon are correlated with the price of electricity, the profit margin of a gas-fired generator can be less risky than either its costs or its revenues, considered in isolation. The profit margin of a nuclear generator may be much more risky than that of the gas-fired station, since its costs will not be as correlated with the price of electricity.

This paper asks how the choice between a carbon tax and an emissions permit scheme (based on auctions) affects the risk of investing in gas-fired or in nuclear generation. With a carbon tax, the variable cost of gas-fired generation is raised by a fixed amount, which will normally feed through into electricity prices. With carbon permits, the price of emissions depends upon market conditions. In the specific case of the electricity industry, the key factor is the relationship between coal and gas prices. Taking a slightly Euro- and electricity-centric view, the price of carbon will be set so that the electricity industry within the EU reduces its emissions to the level that is needed in order that the demand for emissions permits equals the supply. Assuming that fuel prices are such that coal-fired generation would be cheaper than gas-fired, in the absence of the ETS, this implies that the price of carbon has to rise until enough coal-fired stations have been displaced by gas generation. The price of carbon is then at the level which equalises the cost of some gas-fired and coal-fired stations (Newbery, 2006). A rise in the price of gas will tend to raise the price of permits. The higher gas price thus has both a direct effect on the price of electricity, and an indirect effect via the cost of emissions.

The aim of this paper is to discover whether this magnified relationship between gas and power prices has a significant impact on the relative risk of different kinds of generator, compared to the alternative of a carbon tax that was not linked to the level of fuel prices. The paper follows the same basic approach as the papers already cited, but uses a more detailed model of the relationship between input prices and the wholesale price of electricity, that of Evans and Green (2005). This allows us to calculate the expected profitability of each type of power station, taking into account the way that its operating pattern will depend upon its variable cost relative to other stations.

The next section briefly discusses the background to the Emissions Trading Scheme and the economics of a liberalised electricity industry. Section 3 outlines the supply function model used to determine the relationships between fuel and carbon prices and generator profits. Section 4 presents the data used, drawn from the DTI Energy Review (2006) and Supergen FutureNET scenarios (Elders *et al*., 2006). Section 5 presents the results for single plants, and section 6 discusses optimal portfolios containing a mix of plants. Section 7 concludes.

2. THE ELECTRICITY INDUSTRY AND THE EMISSIONS TRADING SCHEME

Traditionally, the electricity industry has largely been vertically integrated. Large companies that combined generation and transmission might sell power to smaller distribution utilities, but this was usually done via contracts or tariffs, rather than with any kind of market mechanism. Following Chile, England and Wales, and Norway, many countries have now adopted wholesale markets for electricity, with competition between generators. While the details vary across countries, the key elements are that generation has been split from transmission, that entry into generation is largely deregulated, and that generators compete to sell their output, through a centralised market, bilateral contracts, or both.

This has changed the way in which companies need to think about investment. Traditionally, investment plans were made with the objective of minimising the expected cost of meeting the forecast level of demand. There could be a trade-off between capital costs and fuel costs – plant that was expected to run for most of the time could incur high capital costs in return for lower fuel costs (the standard example being nuclear power) while still being competitive against plant with lower capital costs but high fuel costs. For plant that was not expected to run for much of the time, it would not be worth incurring the high capital costs. Most newly-built plant *was* expected to run nearly continuously on base load, however, and so valid cost comparisons could be made on the basis of the expected cost per kWh generated at a standardised, high, load factor. The option with the lowest expected levelised cost would normally be the front-runner for investment. The classical investment appraisal appeared to take little notice of risk or uncertainty.

In a market-based system, companies invest to earn profits. Their first criterion will not be to minimise the expected cost of meeting demand, but to maximise the expected difference between their costs and their revenues. In a fully competitive market, in which the company was not able to influence the prices a plant received, this would come down to minimising its expected cost, as with the pre-liberalisation approach. This equivalence between the outcome of a perfect competitive market and a perfect social planner is, after all, one of the arguments for liberalisation.

Companies in the market-based system have a second criterion beside the level of profits, however. Typically, they wish to avoid taking on excessive risk. One approach for this is to sign long-term contracts for fuel inputs and electricity sales at the start of construction, locking in the plant's selling margin. This depends on finding counter-parties willing to take on the risk of these contracts. In Finland, a group of energy-intensive industrial customers has contracted to take the output from a new-build nuclear plant, which should give them stable power costs throughout their own long-lived investment cycle. A company selling to captive customers would also be a natural counter-party for the electricity contract, but liberalisation may have deterred electricity retailers who might lose their customers from taking on such commitments (Newbery, 2002). In the absence of a contract to lock in the risk, companies will have to trade off risk and expected profit when deciding on investments.

In a liberalised market, the wholesale price of power should be related to its marginal cost. The more competitive the market, the closer prices should be to marginal cost. At times of peak demand, this marginal cost should be understood to include the cost of rationing demand if this is necessary to match it with the available capacity. At other times, marginal cost consists of fuel and variable operations and maintenance costs, with fuel costs dominating.

The marginal cost of the system naturally depends upon which plants are operating, since the most expensive of these will determine the system's marginal cost. The approach

used in the traditional system was essentially to stack the plants in a merit order of increasing marginal cost, and to call on the lowest-cost plants first.[1] Nuclear plants were expected to run on base load, while high (marginal-) cost open cycle gas turbines would only operate at times of high demand. These stations suffered from both a low thermal efficiency and a fuel that was usually expensive. Other stations might trade off a higher fuel price against a greater thermal efficiency – this was the situation with Combined Cycle Gas Turbines in the UK for much of the 1990s, which paid more per kWh of fuel than coal-fired plants did, but needed less fuel per kWh generated, giving them a lower marginal cost. When fuel costs change, however, the merit order will change.

In the European Union, a new component has recently been added to generators' marginal costs, with the introduction of the Emissions Trading Scheme. Adopted in response to the Kyoto Treaty on greenhouse gas emissions, the ETS requires all large combustion plants in the EU to surrender an emissions permit for each tonne of carbon dioxide that they emit, or pay a penalty. These permits are issued by Member States, following national allocation plans that have to be approved by the European Commission. The intention is that the plans will be consistent with each country's commitments under the Kyoto Treaty and the EU's burden-sharing agreement, which provides for differential reductions relative to each country's 1990 baseline, depending on its individual circumstances. If "business as usual" would involve more emissions than are compatible with these commitments, then there will not be enough permits to go round. Their price will rise, and this should signal a need for companies to change their behaviour.

The first phase of the scheme runs from 2005 to 2007, with a second phase from 2008 to 2012, covering the Kyoto commitment period. So far, permit prices do appear to be passed through into electricity prices. In 2005, prices in Germany and the Netherlands rose by between 60% and 117% of carbon costs (Sijm *et al*, 2006). This is in line with the predictions of economists, but not perhaps of politicians who expected that if most permits were granted free of charge, keeping companies' average costs unchanged, then they would not change their prices (Gabriel, 2006).

The free allocation can have important effects on companies' investment behaviour, however (Green, 2005). If permits are given away to plants that continue to operate or are newly built, but not to those that close, this can be seen as a subsidy towards their fixed costs, coupled with a tax on their marginal costs. The tax is passed through to prices, while the subsidy gives an incentive to keep more capacity open. To the extent that peak prices depend on the margin of capacity, this effect can offset some of the scheme's impact on marginal costs, limiting the overall effect on average prices in the long run. The way in which permits are given away can also affect the choice of technology. If the allocation is linked to past emissions, or to the technology chosen by a new plant, then more polluting plants will tend to get a more valuable allocation of permits. They will also need more permits, of course, but the greater allocations reduce the incentive to switch investment decisions towards cleaner technologies, and closure decisions towards more polluting ones. A technology-neutral allocation (which should even include giving away permits to nuclear stations) would maximise the scheme's impact on the choice of technology.

If giving away carbon emissions permits to nuclear plants seems politically unattractive, an alternative would be to auction the permits. Since the combination of free

[1] In practice, a more complicated optimisation is required to take account of constraints such as limits to the rate at which plants can change their output, and the cost of starting up a plant to meet a peak in demand, which tends to raise the true marginal cost at the peak well above the "steady state" variable cost of meeting that level of demand if it were to persist for a long time. Similarly, it is expensive to restart a plant after switching it off when demand is low, and the benefit of avoiding this reduces the marginal cost at times of low demand. These complications are not important for the purposes of this paper.

allocation and permit prices feeding through to power prices gives electricity generators a windfall profit, auctioning the permits would claw this profit back to governments. The UK government has announced that it will auction 7% of its permits in phase II of the ETS – the rules for the scheme set an upper limit of 10% for this phase. It remains to be seen whether this limit will be relaxed in future phases (and indeed, whether there will be any future phases).

In a world of certainty, auctioning permits starts to look very like an alternative method of dealing with externalities, a Pigovian tax. If the price of the permit could be predicted, the same amount could be imposed as a tax, with very similar economic effects. It is quite possible, however, that the costs of administering the tax would be lower than those of trading permits. In an uncertain world, the two schemes will have different effects, as one fixes the quantity of emissions, while the other fixes the amount that companies are prepared to spend to avoid it. Weitzman (1974) shows that if the marginal damage from pollution and the marginal cost of avoiding it are uncertain, the optimal economic choice between the two systems depends on the relative slopes of the functions relating marginal damage and cost to the level of pollution. In the case of global warming, the marginal damage is believed to be quite insensitive to the level of emissions in any one year, implying that a tax would be a more suitable economic instrument. The politics of rent-seeking favour a permit system, however, as interest groups can argue for the free allocation of valuable permits, as has happened with the ETS.

What determines the value of permits? The answer is supply and demand, or rather perceived supply and demand. When the scheme started, data on carbon emissions in the EU was very incomplete, and the release of the first comprehensive data in May 2006 caused a significant fall in permit prices. Market participants realised that emissions had been lower, relative to the supply of permits, than they had believed, implying that the price needed to align the two was also lower. It is unlikely that there will be significant investment effects during phase I of the ETS, given the short timescales involved, and balancing emissions with the supply of permits depends upon operating decisions. If the price of electricity rises, demand for it should fall, reducing emissions, but the key operating decision will be between gas- and coal-fired power stations. Gas contains carbon and hydrogen, while almost all the energy content of coal comes from carbon, implying that burning coal will produce more CO2 per unit of energy than burning gas. Furthermore, many gas-fired stations in Europe are combined cycle gas turbines with higher thermal efficiencies (lower heat rates) than coal-fired plants. This implies that switching from coal-fired to gas-fired stations will typically reduce emissions. If the price of gas is low, gas-fired plant will naturally be above coal plants in the merit order, and emissions will be low. If the price of gas is high, however, coal will be favoured. If this is incompatible with the number of permits available, then the price of permits should rise in order to reduce coal's cost advantage and give generators an incentive to burn more gas. We would then expect the price of permits to be positively correlated with the price of gas, and negatively correlated with that of coal.

These correlations, and the correlations between fuel and electricity prices, are at the heart of what follows. If generators are worried about profit risks, then they will have an incentive to choose technologies with costs that are correlated to power prices. The costs of gas-fired stations have this property, while those of nuclear stations do not. If carbon prices are also correlated with gas prices, then this will increase the volatility of both power prices and gas-fired generators' costs, while they remain correlated. Nuclear stations are not exposed to carbon prices and will be relatively more risky. The next sections of the paper set out the model that is used to test this theory.

3. A SUPPLY FUNCTION MODEL

This paper uses a supply function model to predict generators' profits, given input costs and the level of demand. Klemperer and Meyer (1989) introduced the supply function equilibrium, while Green and Newbery (1992) applied it to the British electricity market. As argued in that paper, the supply function equilibrium is a close approximation to the workings of the Pool, in which companies effectively had to submit offers of prices and quantities (from each of their many power stations) that would hold throughout the following day.[2] These offers can be represented by a supply function, and the equilibrium price and output in each period are given by the intersection of the aggregated supply function with the market demand curve. Demand varies over time, which is mathematically equivalent to the stochastic variation considered by Klemperer and Meyer.

Formally, demand is denoted by $D(p,t)$. Assume that $dD/dp < 0$, and that $d^2D/dp^2 \leq 0$. There are n generators, which compete by submitting supply functions $(q(p): R \rightarrow R, i = 1...n)$ which state the amount they would be willing to produce (q) at any price (p). These functions must be non-decreasing in p - the Pool's rules ensure this by ranking plants in order of increasing bids. The price at each time is determined by a market-clearing condition. The total output supplied at the market-clearing price must just equal the demand with that price at that time:[3]

$$D(p^*(t),t) = \sum_i q_i(p^*(t)) \tag{1}$$

An equilibrium consists of a set of supply functions, one for each firm, such that each firm is maximising its profits, given the supply functions of the other firms, at every time. We can write each firm's profits π_i (revenues, less the cost $(C(q))$ of production) at each time as a function of price, assuming that it produces the residual demand (that is, total demand less the other firms' supply at that price) in order to meet the market-clearing condition:

$$\pi_i(p,t) = p\left(D(p,t) - \sum_{j \neq i} q_j(p)\right) - C_i\left(D(p,t) - \sum_{j \neq i} q_j(p)\right) \tag{2}$$

This profit function can be differentiated with respect to price:

$$\frac{\partial \pi_i(t)}{\partial p} = D(p,t) - \sum_{j \neq i} q_j(p) + \left(p - C_i'\left(D(p,t) - \sum_{j \neq i} q_j(p)\right)\right)\left(\frac{\partial D(p,t)}{\partial p} - \sum_{j \neq i} \frac{\partial q_j}{\partial p}\right) \tag{3}$$

Setting this derivative to zero gives the profit-maximising price level at a particular time, and it also gives the profit-maximising output (i.e., the residual demand) at that price level. Assume that $\partial^2 D/\partial p \partial t = 0$, and then this price cannot be optimal for a different level of demand. The (price, quantity) pair will then form a point on the profit-maximising supply function. We can manipulate the first-order condition to give a differential equation for the firm's supply function:

[2] Companies could vary their available capacity each half-hour, but very rarely did so in a strategic manner (as opposed to not bothering to staff a peaking plant which would not run overnight, for example).

[3] For completeness, we assume that if there is no price which solves this condition, the price will be zero.

$$q_i(p) = \left(p - C_i'\left(q_i(p)\right)\right)\left(-\frac{\partial D}{\partial p} + \sum_{j \neq i}\frac{\partial q_j}{\partial p}\right) \tag{4}$$

A supply function equilibrium consists of a set of solutions to equation (4), one for each firm, such that the resulting functions are sloping upwards for every price that might be obtained from the intersection of a demand curve and the aggregate supply function. If the firms are symmetric, there is a wide range of potential supply functions, although it narrows as the variation in demand increases.

Evans and Green (2005) were able to simulate electricity prices from April 1997 to March 2005, using a symmetric approximation of the industry supply function. They modelled the industry as if it contained $\hat{n}$ symmetric strategic generators, together with a competitive fringe. In each month, $\hat{n}$ was the inverse of the Herfindahl index, calculated using the capacity of the strategic generators. This gives a straightforward differential equation:

$$q_i(p) = \left(p - C_i'\left(q_i(p)\right)\right)\left(-\frac{\partial D}{\partial p} + (\hat{n}-1)\frac{\partial q_i}{\partial p}\right) \tag{5}$$

Since only the industry supply function is required to predict prices and (overall) operating patterns, it is not necessary for $\hat{n}$ to be an integer. Following Evans and Green, this paper uses the highest-priced supply function that includes all of the firms' capacity – it starts from the price that the firms would charge if they were selling their full capacity in a Cournot equilibrium. This industry supply function is then combined with the supply curve of nuclear stations bidding at their near-zero marginal cost to give an overall supply curve, shown in figure 1. The predicted equilibrium price for a given period is then given by the intersection of this supply curve with the demand curve for that period. We can also find the marginal cost at this point, and hence infer which stations are generating.[4] This implicitly assumes that the industry truly *is* symmetric, since a smaller generator will produce more, relative to its capacity, than a larger one, and so marginal costs would not be equalised across asymmetric generators. For the purposes of this paper, this simplification is unlikely to create important effects.

4. CALIBRATION

Much of the data for this paper comes from the UK government's recently-published Energy Review (DTI, 2006, Annex B). Demand in 2020 is expected to be around 400 TWh in the DTI's central case. The pattern of demand within the year is assumed to be the same as in 2003-5, scaled up by 28% to give the desired total.

The review includes information on the costs of a range of generating technologies, and predictions of fossil fuel prices, expressed as central, high and low cases. The variable cost of a power station is equal to the price of its fuel and carbon (permit or tax) per kWh of fuel, divided by its thermal efficiency, plus its variable operations and maintenance costs. The review also gives information on fixed operations and maintenance costs, capital costs (per kW) and the expected life of the plants. The capital costs were converted to annual

[4] With continuous marginal costs, there is a specific marginal point on the cost function, and all capacity with marginal costs below this level would be operating. This paper uses a step function, and so it will typically be the case that not all of the marginal tranche of capacity will be operating – we assume that each plant in the tranche has the same chance of running. When calculating the variability of profits, however, we treat a 50% chance of running as a deterministic period of operation at 50% of capacity.

figures with a discount rate of 10%, the rate used by the DTI.[5] Where the DTI gives two figures, the mid-point is used.

The three main technologies considered in the industry supply function are coal-fired stations, combined cycle gas turbines, and nuclear stations. A small amount of oil-fired plant, and prototype stations with carbon capture and storage, was also included in the model. These higher-cost plants have the effect of raising the industry's marginal cost at the far end of its supply function, and hence raising the price received by all the other plants. The amounts of coal, gas, and nuclear plants were taken from one of the Supergen FutureNET scenarios for 2020 (Elders *et al.*, 2006), "supportive regulation". This scenario has a slightly greater demand for electricity than the DTI's base case, however, and so this paper assumes slightly less CCGT capacity than the scenario. We assume a Herfindahl index of 1/6, equivalent to six equal-sized firms, when deriving the supply function for the industry's non-nuclear capacity. Nuclear stations are assumed to run on base load and not to affect the market price.

The prices of gas, coal, and oil are stochastic, with means equal to the DTI's central predictions for 2020. All three have lognormal distributions. The price of oil is taken to be the underlying random variable, while the prices of gas and coal are linked to it.

$$\begin{aligned} p_{oil} &= \exp\left(0.33\eta_{oil} + 2.7175\right) && ; \quad \eta_{oil} \in N(0,1) \\ p_{coal} &= \exp\left(0.38\left(0.5\eta_{oil} + 0.5\eta_{coal}\right) + 1.35\right) && ; \quad \eta_{coal} \in N(0,1) \\ p_{gas} &= \exp\left(0.36\left(0.7\eta_{oil} + 0.3\eta_{gas}\right) + 2.4875\right) && ; \quad \eta_{gas} \in N(0,1) \end{aligned} \qquad (6)$$

The distributions are calibrated so that each fuel price has a mean equal to the DTI's Central Case (Favouring Coal) for 2020. The means plus or minus two standard deviations are approximately equal to the DTI's High and Low cases for each fuel. (The DTI's figures are not quite consistent with a lognormal distribution, and so this match cannot be exact). Table 1 shows the mean and standard deviations for each fuel, together with their correlation coefficients. The bottom part of the table gives the correlations between real annual average fuel prices using data from the BP digest of energy statistics for Brent crude, and coal and oil delivered to North-West Europe. This assumes that as the UK imports an increasing proportion of its gas, its prices will converge on those in North-West Europe. Our main simulations match the higher correlations between fuel prices seen in the more recent past, but a second set, based on the lower correlations, give very similar results in both quantitative and qualitative terms.

Table 1: Fuel prices (£/MWh)

		Oil	Gas	Coal
	Mean	15.99	12.48	3.99
	Standard Deviation	5.40	3.39	1.09
	Correlation with	Oil	0.89	0.70
		Gas		0.63
Correlation in BP world energy tables	1996 to 2005	Oil	0.89	0.70
		Gas		0.62
	1987 to 2005	Oil	0.86	0.50
		Gas		0.48

[5] A much lower rate, of 2.2%, was used to discount the costs of nuclear decommissioning and waste reprocessing.

In the model runs for the ETS, the price of carbon is set to vary randomly around the level that would equalise the cost of generation from a coal-fired and a gas-fired station. As discussed above, this is based on the view that switching between coal and gas will be necessary to keep emissions down to the level required by the ETS, and that the price of carbon will vary around the level that would keep plants of both types on the margin. The relationship between the costs of two power stations is given by:

$$\frac{0.3395 p_{car} + p_{coal}}{0.35} = \frac{0.19635 p_{car} + p_{gas}}{0.53} \qquad (7)$$

where p_{car} is the price of permits per tonne of carbon dioxide, p_{coal} is the price of coal per MWh, and p_{gas} is the price of gas per MWh. The coefficients on the fuel prices are the amounts of carbon dioxide per MWh of fuel, while the denominators are the assumed thermal efficiencies of the plants that are marginal in terms of fuel switching to meet the emissions limit. A random normal variable was then added, to reflect uncertainties over how far generators would have to switch between coal and gas to meet the emissions limit, for example. This gives us an equation for the price per tonne of carbon dioxide:

$$p_{car} = 3.15\, p_{gas} - 4.77\, p_{coal} + \text{N}(0,2) \qquad (8)$$

This price was not, of course, allowed to become negative. This gave a mean carbon price of £20.18 per tonne, and this was the level imposed as a carbon tax. This is close to the level (£17/tonne) that the DTI used in its Central case in the Energy Review. For a sensitivity analysis, a carbon price with the same mean but half the link to fuel prices was also used:

$$p_{car} = 10.11 + 1.57\, p_{gas} - 2.38\, p_{coal} + \text{N}(0,1) \qquad (9)$$

The paper so far has not discussed the role of the Clean Development Mechanism (CDM) in the ETS. Instead of reducing their own emissions or buying ETS permits, companies can pay for emissions reduction measures in developing countries, and surrender the certified emissions reduction units obtained for these. By December 2006, 435 schemes expected to save 106 million tonnes of carbon dioxide (equivalent) a year had been registered, with roughly twice as many in the pipeline and a predicted saving of 1.5 billion tonnes by 2012. If the CDM provides a sufficiently elastic supply of emissions reductions at a sufficiently low cost, then this will cap the price of ETS permits. If gas is cheap and coal expensive, then fuel switching will be the preferred option and will keep the price of permits low. If the price of gas rises, CDM schemes could be used instead. The combination would resemble a hybrid tax-permit scheme which would behave like a trading scheme when fuel setting implied a low permit price, and a tax when the (assumed) elastic supply of CDM credits provided a lower-cost option.

The techniques used in this paper could easily be applied to analyse such a hybrid scheme. The aim of this paper is to show the differences between a tax and a trading scheme, and for this, it is best to work with "pure" versions. If the supply of CDM credits is not sufficiently elastic, they will act to shift the amount of fuel switching (or other measures) required in Europe, but will not be able to act as a price cap. For both these reasons, CDM credits are not explicitly included in the model of this paper.

For each scenario run, one industry-level supply curve was generated for winter and one for summer. For winter, 90% of plant was assumed to be available, with 80% availability in summer. Since the price of gas is seasonal, but the degree of variation varies from year to

year, the price of gas in winter was raised above the annual average obtained earlier by a random amount, uniformly distributed between zero and 12% (in line with past patterns). The summer price was reduced by the same amount, preserving the annual (time-weighted) average. Twenty-six demand curves for representative half-hours were used in each season, based on wholesale market price-quantity pairs from 2003-5 (with the quantity demanded scaled up as described above), taking the lowest demand level from April to September, the 4th percentile, 8th percentile and so on. Renewable generation was subtracted from the gross demand to give the net demand that must be met by large stations. This output is stochastic, with a mean load factor of 33%, and a standard deviation (for nationwide output) of 5%. Once again, stations had an availability of 90% in winter and 80% in summer, giving slightly higher generation levels in the winter months.

The intersections of the net demand curves and the industry supply curve were found, and the equilibrium prices calculated. This allowed us to predict the output from each type of plant in each of the representative half-hours, and the margin between their revenues and variable costs. An individual station would only earn this margin if it was actually available, however, with a probability of 0.9 or 0.8 (depending on the season). If not available, it would earn no margin in that half-hour. The station's expected profits over the year were equal to the 52 margin figures (equally weighted), scaled up for the total number of half-hours in the year, less its fixed costs.

5. RESULTS

The model was simulated 50,000 times for carbon trading, and 50,000 times for a carbon tax. The key results are shown in table 2, and figures 2 and 3. In figure 2, it is clear that with carbon trading, gas-fired stations have the least variable profits, followed by coal-fired stations, while the profits of a nuclear station are dispersed over a wide range that includes a nearly one in four chance of making an annual loss. Comparing this with figure 3, when a carbon tax is imposed instead, the profits of coal and gas-fired plants become more variable, and those of nuclear plants less variable. Nuclear profits are still clearly more dispersed than those of gas-fired stations, and table 2 confirms that they have a greater standard deviation than those of coal-fired plants.

Table 2: The impact of carbon policy

	Profits (£ per kW-year)			Price (Annual average)
Carbon trading	CCGT	Coal	Nuclear	£/MWh
Mean	36.76	9.25	32.89	42.22
Standard deviation	8.41	13.88	64.40	8.58
Carbon tax				
Mean	38.35	10.76	34.54	42.43
Standard deviation	11.55	29.73	38.25	5.04

In this model, new gas-fired stations are slightly more profitable on average than nuclear stations, which are both at least three times as profitable as new pulverised fuel coal plants (without carbon capture). Compared to the variation between plant types, carbon policy has little impact on mean profits, which is by design, since the carbon tax is imposed at the mean level of the carbon price in the trading simulations. There is a slight increase in the average electricity price, and in mean profits, with a carbon tax. The observed increase in the average

price, of about £0.2/MWh, would imply an extra profit of £1.50/kW for a plant that ran 7,500 hours a year with unchanged costs.

The increase in prices comes from the way in which the carbon policies affect the marginal cost curve. By design, carbon trading will produce a relatively flat marginal cost curve over the industry's fossil-fired plants, since the carbon price adjusts to bring the marginal costs of coal and gas-fired plant close together. A carbon tax does not change when one type of plant has relatively lower fuel costs, and so the marginal cost curve will tend to be steeper. The plants with lower costs, whichever they happen to be, will be at the bottom of the curve, and those with higher costs at the top. The absolute position of the marginal cost curve, however, will be more volatile with a system of carbon trading than with a carbon tax. When the gas price is low, a low carbon price will tend to keep marginal costs low, whereas the impact of a high gas price is magnified by that of a high carbon price. This effect does not occur with a carbon tax. The curves are shown, schematically, in figure 4. The outer solid lines are assumed to be the extremes encountered with carbon trading, whereas the inner solid lines are the extremes produced by a carbon tax. The dashed lines show the averages of these two types of marginal cost curve. By construction, the averages cross half-way along their length – the marginal cost of an "average" plant has not been affected by the choice of policy. The equilibrium prices obtained, however, depend upon where the marginal cost curves are intersected by demand curves. If a demand curve runs to the right of the cost curves' crossing point, it will produce a higher average price with a carbon tax than with carbon trading. In these model runs, a majority of the demand curves used are for these higher demand levels, and so the average price across all demand curves and all simulations is slightly higher with a carbon tax than with carbon trading. Had more of the demand curves been to the left of the crossing point, a carbon tax would have led to a slightly lower average price.

In Finland, a nuclear station is being built with the support of long-term contracts with a group of industrial customers, companies that want to have a stable electricity price over the (very) long term. It is straightforward to assess the impact of similar contracts on the economics of our power plants.[6] The contract used is a base-load forward contract that delivers the same amount of power in each half-hour of the year, at a price equal to the time-weighted average price across all the runs for a given carbon policy. This means that the generators' expected profits are not affected if some of their output is sold at this fixed price, rather than at the variable market price, but the variability of their profits will be affected.

A nuclear generator can indeed reduce its risks by a significant amount if it covers much of its output with contracts. With carbon trading, the generator's profits have a standard deviation of about £64/kW-year, but this falls to a little £7/kW-year if the generator signs a contract that covers its expected output. This would be a contract for 0.85 kW of output per kW of capacity, given that the station expects to be unavailable for 10% of the time in winter and 20% of the time in summer. With a carbon tax, signing the same contract reduces the standard deviation of profits to a little under £7/kW-year.

A gas-fired generator does not reduce its risks by signing a contract to sell its output at a fixed price – in fact, selling 0.85 kW of output in advance raises the standard deviation of its profits from £8/kW-year to £64/kW-year with carbon trading, and from £12/kW-year to £44/kW-year with a carbon tax. The problem is the reverse of that faced by the nuclear generator, since the contract fixes the gas-fired station's revenues while leaving its costs variable. If the station could combine fixed-price contracts for gas purchases and for electricity sales, then its risks would fall.

[6] I would like to thank Matthew Leach for suggesting this addition to the paper.

6. PORTFOLIO EFFECTS

Looking at individual stations, investing in CCGT plant is clearly the most attractive option, with higher predicted profits and a lower standard deviation, whether carbon is traded or taxed. Coal-fired plants appear least profitable under either system, while the returns to nuclear plant have the highest standard deviation. Long-run contracts could reduce the variability of nuclear plants' profits, but will not increase them.

In practice, however, most investment decisions in the electricity industry are made by companies that own a portfolio of plant, and we must consider portfolio effects. Awerbuch (2000) has applied portfolio theory from the finance literature to show that including renewable generators in a portfolio of plant can improve the trade-off between expected generating cost and its variability. The portfolio can improve its balance between return and risk by adding an extra asset, even if that asset is inferior on both counts to the rest of the portfolio, considered in isolation, as long as the correlation between the asset returns is not too high. Twomey (2005) and Roques *et al.* (2006a) show how this theory can be applied to different portfolios of thermal plant.

Figure 5 shows the frontiers of risk and return that can be achieved by constructing a range of optimal portfolios for each policy.[7] At the top right hand end of the frontier, the highest return (expected profit per kW) comes from a portfolio of gas-fired plant alone, but this also has the highest risk, measured in terms of the standard deviation. Adding nuclear stations, contracts, and coal-fired plants allows a company to construct a portfolio with a lower level of risk, although it will have to sacrifice some expected return to do so. Note that the proportion of nuclear plant in the low-risk portfolios is well above the proportion in the system as a whole. This might be achievable for a single company, but if every major company in the industry tried to construct a similar portfolio, the industry's plant mix would change. We could run the simulations again for the new plant mix, check whether the optimal portfolio had changed, and eventually reach a fully consistent outcome.

We do not need to know the generators' preferences between risk and return in order to specify the frontier of optimal portfolios, for it is simply the outer envelope of the many points which could have been plotted in figure 5 – interior points were in fact omitted for clarity. To assess which portfolio a generator might choose, however, we need to know their trade-off between risk and return. A standard model of this trade-off uses a mean-variance utility function, given by:

$$U = \pi^e - \tfrac{1}{2}\,\lambda\,\mathrm{var}(\pi) \qquad (10)$$

where U is the generator's utility, π^e is the expected profit (per kW-year), var(π) is the variance of profits per kW-year, and λ is the coefficient of risk aversion. It is possible to use Grinold's (1996) "grapes from wine" technique to infer the value of λ from stock market data on ex-post returns and their variability. UK investors would appear to have a value of λ equal to 2.2 divided by their wealth, where wealth needs to be measured relative to the units in this model (here, kW of capacity) (Green, 2004). With profits of around £50/kW-year from most of the potentially optimal portfolios, and a discount rate of 10%, this is likely to produce a very low value for λ, of less than 0.005. At this level of risk aversion, an all-gas portfolio

[7] Each portfolio took a weighted average of the profits of three individual plants and of a long-term contract, and the mean and standard deviation of each portfolio's profits were recorded. The optimal portfolios were found by a grid search, using steps of 2 percentage points. The variability of portfolios with a high proportion of a single plant type is slightly over-estimated, since they are based on a single unit with stochastic availability. In reality, such a portfolio would include more than one plant, with uncorrelated availability and slightly less variable profits.

would be optimal, but we will show the way in which the optimal portfolio changes in response to rising risk aversion over a much greater range of values for λ.

Figures 6 and 7 show the optimal portfolios obtained for different degrees of risk aversion under the two policies, assuming that the companies are unable to sell long-term contracts to hedge the risks faced by nuclear plant. With carbon trading, the optimal proportion of nuclear plant is tiny, and this is edged out by coal plant for high degrees of risk aversion. With a carbon tax, there is a noticeable increase in the optimal proportion of nuclear plant, to nearly 20%, although once again this is replaced by coal-fired stations for high degrees of risk aversion. Note that the horizontal scale of these diagrams is not linear, for more bars are given for low degrees of risk aversion than for the higher levels.

Figures 8 and 9 show the impact of allowing the generators to sign long-term sales contracts at the expected wholesale price. There is a dramatic increase in the optimal proportion of nuclear plant, most of which is covered by contracts. Once again, a carbon tax is associated with a higher level of nuclear plant for most levels of risk aversion. Coal plants form part of the portfolio for the highest levels of risk aversion, but at a lower proportion than if contracts are not used.

7. CONCLUSION

This paper has used simulation methods to predict the mean and variance of the profits earned by different types of generators in a liberalised electricity market with different types of emissions policy. Using numbers appropriate for the UK in 2020, gas-fired stations appear to be both most profitable and to have the lowest variance of profits. This is due to the high correlation between their marginal costs and the price of electricity. It is likely to be self-reinforcing, since these characteristics make gas-fired stations more attractive to investors, and an increase in the proportion of gas-fired generation strengthens the link between the cost of gas-fired stations and the price of power.

Nuclear power stations have costs that are much less volatile than those of gas-fired stations, but their profits are lower and more volatile. If companies are sufficiently risk-averse, including nuclear power stations in a portfolio of plant could be worthwhile, given the negative correlation between nuclear and CCGT profits. The level of risk aversion required for this diversification argument to be effective, however, could be higher than that supported by estimates drawn from the UK stock market. A higher carbon tax, or more restrictive emissions trading scheme, could raise the profits of nuclear power, but will not reduce their volatility while wholesale prices are linked to the cost of gas. Long-term electricity sales contracts would be needed if the hedge that nuclear generation provides for the system's costs is to become a hedge for company profits.

Finally, it is worth pointing out that a carbon tax only reduces the volatility of nuclear profits if generators can be sure that, once set, it will remain constant or move along a pre-announced path. There may be significant political difficulties in adopting a carbon tax in the first place – the ETS was adopted after several years of attempting to impose a European carbon-energy tax. Ensuring that a tax, if adopted, was not continually adjusted by politicians might be an even greater challenge.

References

Awerbuch, S. (2000) "Getting It Right: The real cost impacts of a renewables portfolio standard" *Public Utilities Fortnightly*, Feb 15, 2000, pp. 44-52

Department of Trade and Industry (2006) *The Energy Challenge: Energy Review Report 2006*, Cm 6887, London, The Stationery Office

Elders, I., G. Ault, G. Burt, J. Koehler, R. Tumilty and J. McDonald (2006) *Scenarios of the Electricity Industry in Great Britain in 2020: Networks, Generation and Decarbonisation*, mimeo, University of Strathclyde

Evans, J.E. and R.J. Green (2005) *Why Did British Electricity Prices Fall After 1998?* Department of Economics Working Paper 05-13, University of Birmingham

Gabriel, S. (2006) *Germany remains at the forefront of climate protection*, downloaded from Federal Ministry for the Environment, Nature Conservation and Nuclear Safety web site on 26 July 2006, http://www.bmu.de/english/emissions_trading/current/doc/37401.php

Green, R.J. (2004) *Retail Competition and Electricity Contracts*, CMI Working Paper EP33, University of Cambridge

Green, R.J. (2005) "Electricity and Markets" *Oxford Review of Economic Policy*, vol. 21, no.1, pp. 67-87

Green R.J., and D.M.Newbery (1992) "Competition in the British Electricity Spot Market" *Journal of Political Economy* 100 (5), October, pp 929-53

Grinold, R.C. (1996) "Domestic grapes from imported wine", *Journal of Portfolio Management*, vol 26, special issue (December), pp 29 - 40

Klemperer, P. D. and M.A. Meyer (1989) "Supply Function Equilibria in Oligopoly under Uncertainty," *Econometrica*, vol. 57 no. 6, pp 1243-1277

Newbery, D.M. (2002) "Problems of liberalising the electricity industry" *European Economic Review*, vol. 46, pp 919-927

Newbery, D.M. (2006) *Climate change policy and its effect on market power in the gas market* Electricity Policy Research Group working paper EP06/06, University of Cambridge

Roques, F.A., W. J. Nuttall and D.M. Newbery (2006a) *Fuel-mix diversification incentives in liberalized electricity markets: a Mean-Variance Portfolio Theory Approach*, mimeo, University of Cambridge

Roques, F.A., W. J. Nuttall and D.M. Newbery (2006b) *Using Probabilistic Analysis to Value Power Generation Investments Under Uncertainty*, Electricity Policy Research Group working paper EP06/19, University of Cambridge

Roques, F.A., W. J. Nuttall, D.M. Newbery, R. de Neufville and S. Connors (2006c) "Nuclear Power: a Hedge against Uncertain Gas and Carbon Prices?" *The Energy Journal* vol. 27 no. 4, pp 1-23

Sijm, J., K. Neuhoff and Y. Chen (2006) "CO_2 cost pass-through and windfall profits in the power sector", *Climate Policy*, vol. 6 no. 1, pp.47-70

Twomey, P. (2005) *The non-CO2 benefits of renewables*, presentation to the Electricity Policy Research Group Spring Seminar, May 2005, available from http://www.electricitypolicy.org.uk/events/spring05/programme.html

Weitzman, M.L. (1974) "Prices versus quantities" *Review of Economic Studies,* vol. 41 no.4, pp. 477-91

Figure 1: Industry supply function - DTI Base Case

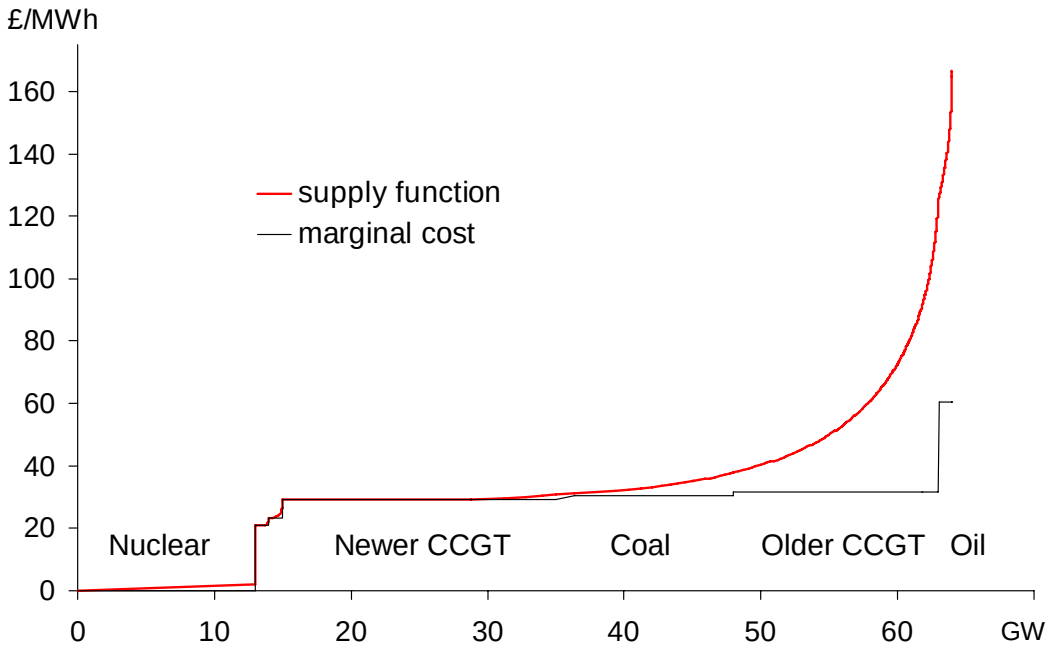

Figure 2: Profits with carbon emissions permits

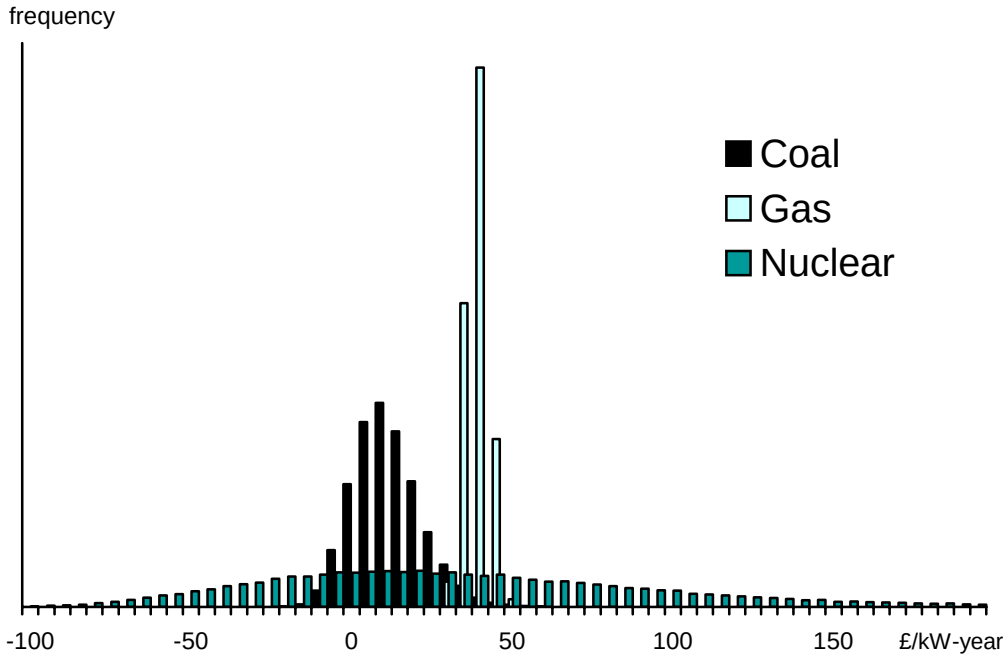

Figure 3: Profits with a carbon tax

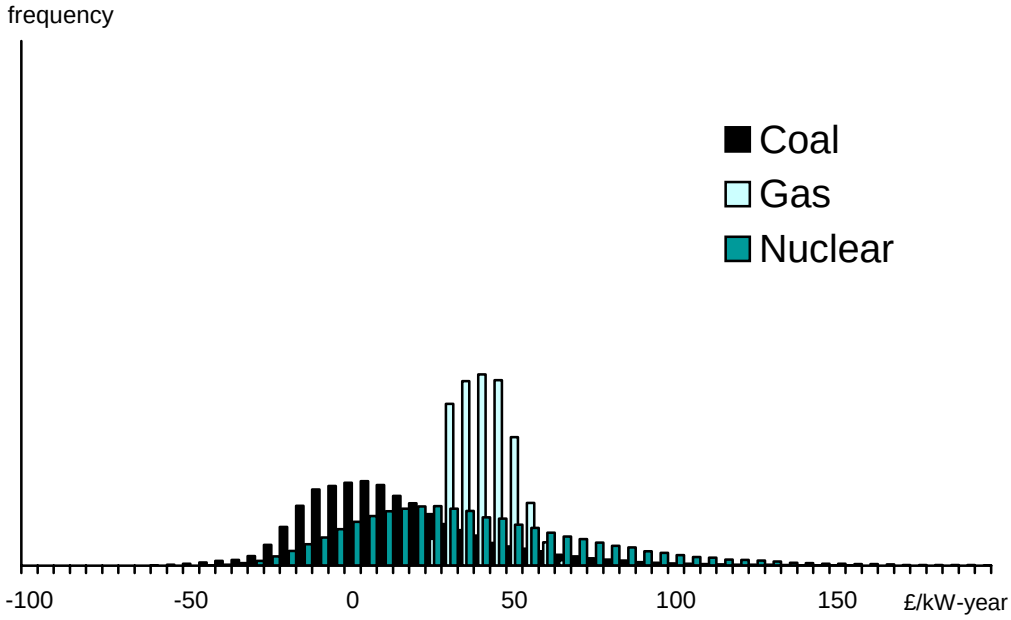

Figure 4 – Why a carbon tax gives higher average prices

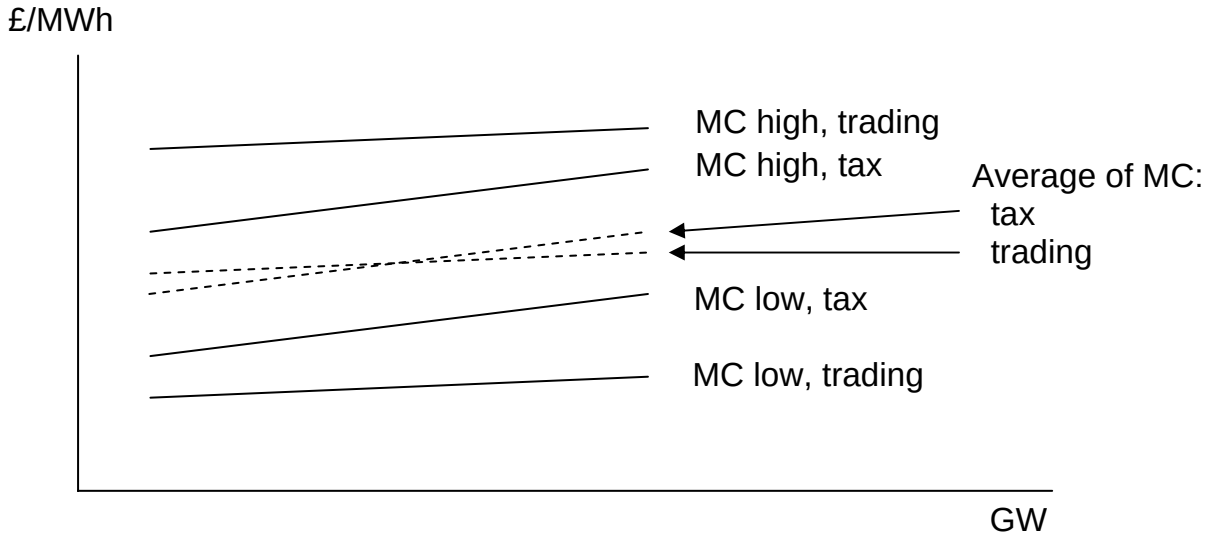

Figure 5: Return and risk from optimal portfolios

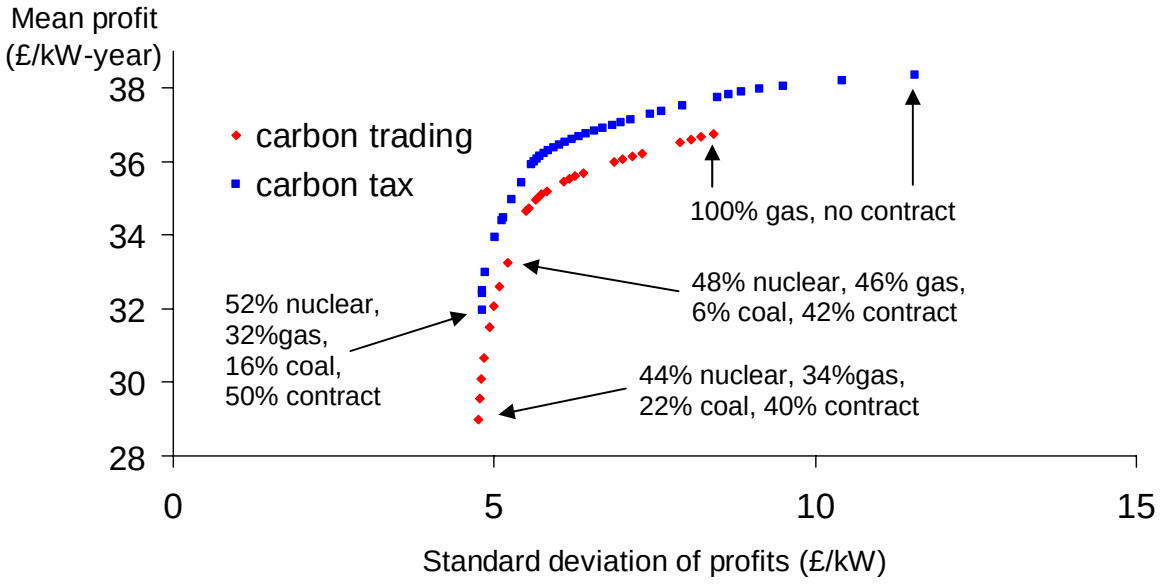

Figure 6: Optimal portfolios with emissions trading

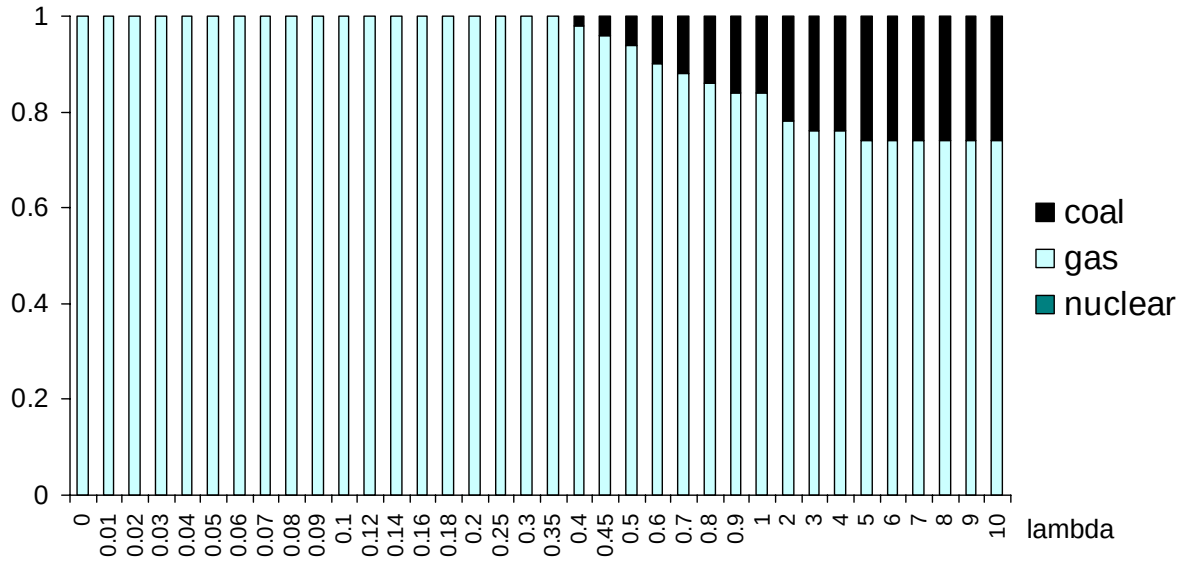

Figure 7: Optimal portfolios with a carbon tax

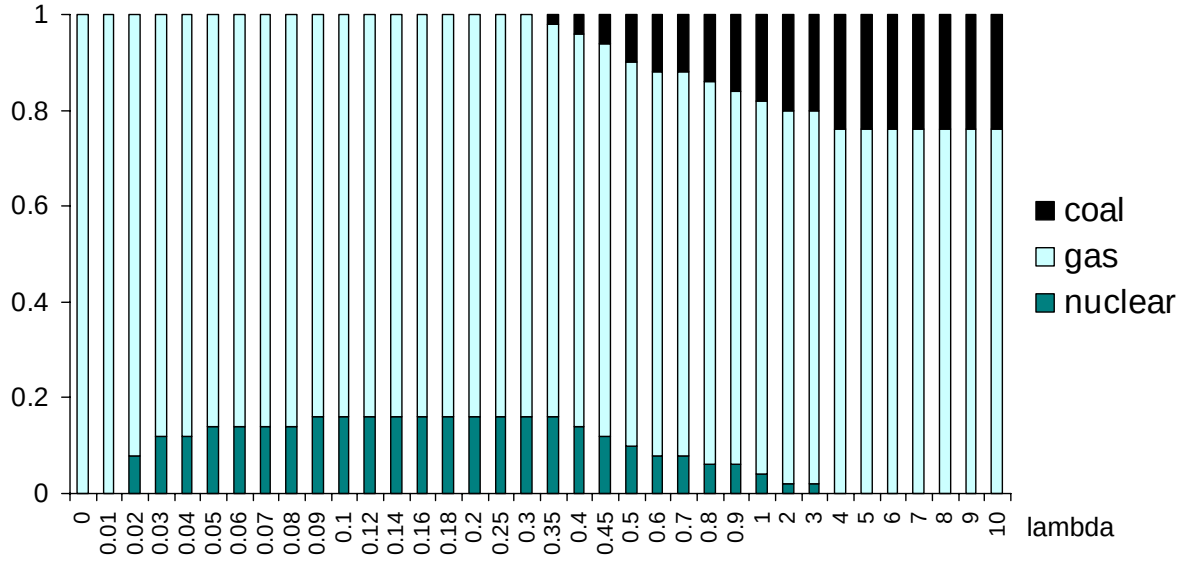

Figure 8: Optimal portfolios with emissions trading and long-term contracts

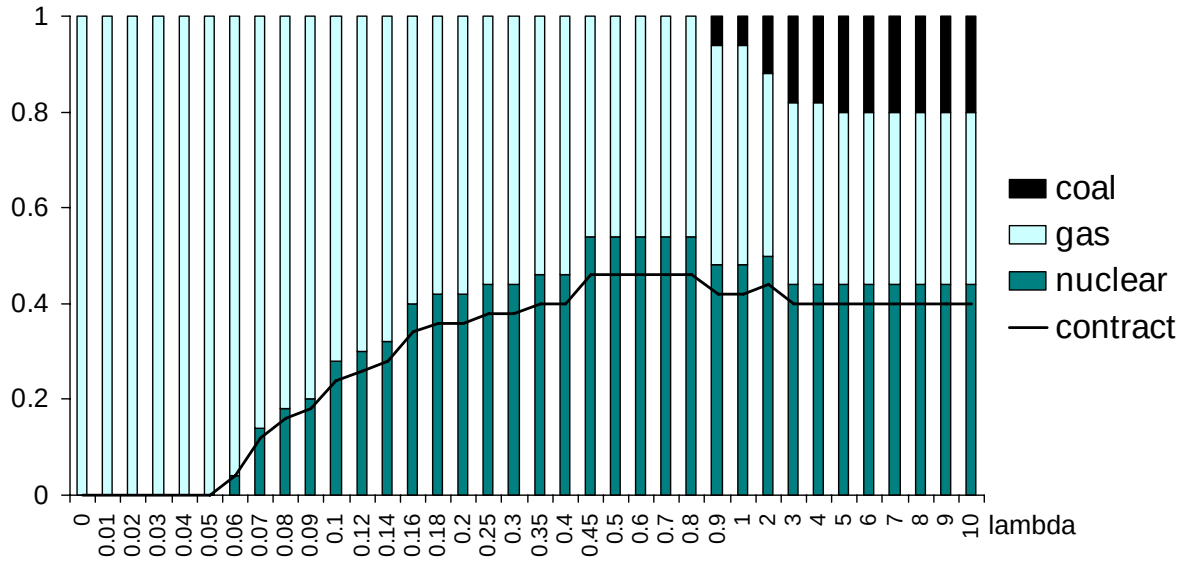
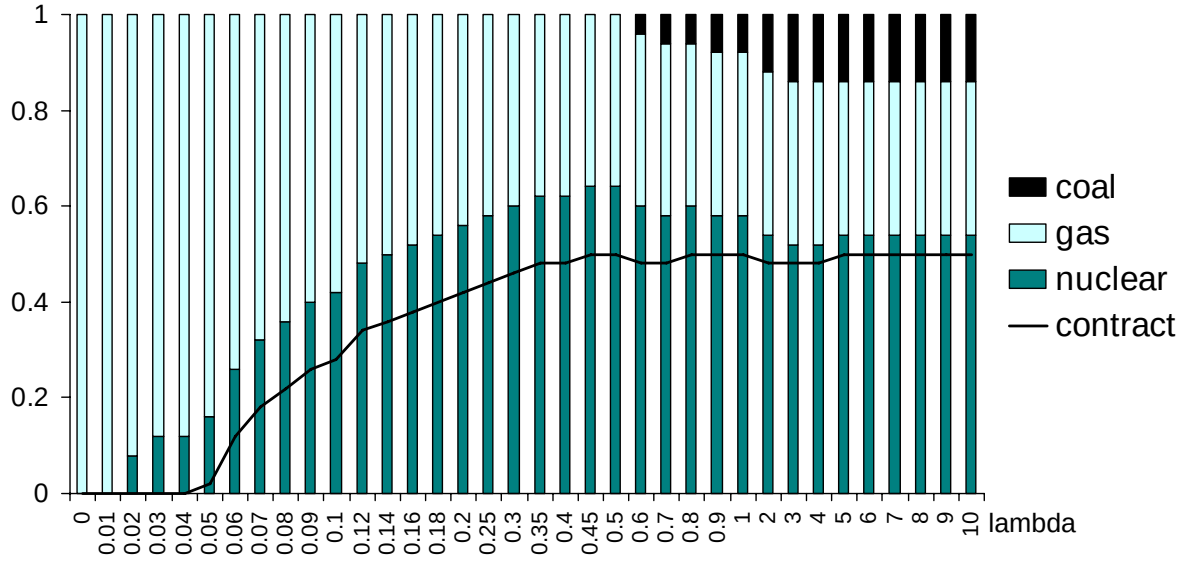
Figure 9: Optimal portfolios with a carbon tax and long-term contracts