



Centre for Energy Policy and Economics
Swiss Federal Institutes of Technology

An Options Pricing Approach for CO₂ Allowances in the EU ETS

Beat Hintermann

CEPE Working Paper No. 64
June 2009

An Options Pricing Approach for CO₂ Allowances in the EU ETS

Beat Hintermann^a

Abstract

If firms are unable to fully control their emissions, the cap in a permit market may be exceeded. Using stochastic aggregate emissions as the underlying I derive an options pricing formula that expresses the permit price as a function of the penalty for noncompliance and the probability of a binding cap. I apply my model to the EU ETS, where rapid market setup made it difficult for firms to adjust their production technology in time for phase 1. The model fits the data well, implying that the permit price was driven by firms hedging against stochastic emissions rather than marginal abatement costs.

Keywords: Permit markets, air pollution, CO₂, climate change, options pricing, EU ETS.

JEL classification: G12, G14, G18, Q52, Q53, Q54

a: ETH Zurich, Centre for Energy Policy and Economics (CEPE); bhintermann@ethz.ch

Introduction

The centerpiece of emissions permit market theory is that firms equate their marginal abatement costs to the permit price at all times. If a firm finds abatement of an additional unit of emissions to be cheaper than the permit price, it will make a profit from abating and either buy one fewer or sell one more permit on the market. Conversely, if purchasing a permit is cheaper than abating another unit of emissions, the firm will use the permit market to reach compliance. The efficient outcome of this arbitrage game is that all firms abate to the point where their marginal abatement costs equal the permit price, and that the permit market clears.

However, real-world permit markets may not clear for two reasons: First, abatement may not be feasible for the involved firms in the short run without cutting output. And second, cutting output can be either economically or legally infeasible, for example in the case of electricity markets where output has to be matched with supply at all times in order for the grid not to crash. It is thus possible that aggregate emissions are stochastic and turn out to be either significantly below or above the cap. Most permit markets to date impose a penalty for noncompliance per unit of emission for which no permit can be surrendered.

The price path of emission allowances during the first few years of European Union Emissions Trading Scheme (EU ETS) has puzzled market participants and economists alike. A series of recent studies analyze the first phase of the market under the implicit assumption of a clearing permit market but find little evidence that the permit price was indeed driven by marginal abatement costs (Alberola et al., 2008; Bunn and Fezzi, 2008; Mansanet-Bataller et al., 2007; Rickels et al., 2007; Paolella and Taschini, 2006).

I set up a model in which firms are unable to effectively control their emissions and buy permits in order to hedge against the possibility of having to pay a penalty. The permit price thus becomes a binary option (a.k.a. “cash or nothing” option) that is a function of the probability of a binding cap and the penalty for noncompliance, but not of marginal abatement costs. I estimate the free parameters from the options formula with daily data from the first

phase of the EU ETS and proxy for daily emissions using daily generation of electricity from fossil fuels. The model fits the data well, at least better than abatement-based models. The fit is greatly increased if a parameter multiplying the variance of future electricity generation is introduced into the model. My findings imply that allowances in the first phase of the EU ETS were driven by firms hedging against stochastic emissions, rather than by marginal abatement costs. This has important efficiency implications, because the equality of marginal abatement costs and permit price is a necessary condition to achieve a given emissions cap at least cost, which is the explicit goal of every permit market.

My paper is closest in spirit to (Chesney and Taschini, 2008), who derive an options model for emissions permits under the assumption of no abatement. The main difference between our papers is that they simulate an underlying pollution process and based on this compute an options price, whereas I take the allowance price and the pollution process as given and estimate a set of free parameters using an options pricing formula.

Section 2 gives some background about permit pricing and the EU ETS. In Section 3 I derive an options pricing formula for EU ETS allowances as a function of emissions, the cap, the penalty for noncompliance and a set of free parameters. This formula contains the mean and variance of expected future emissions, which I derive in Section 4 as a function of exogenous stochastic processes. Section 5 presents empirical estimates for these underlying processes and the free parameters in the options pricing formula. Section 6 concludes.

Background

2.1 Literature

Historically, permit pricing formulas were derived by solving an optimal control problem, originating with Montgomery (1972). This was later extended to the dynamic case (Leiby and

Rubin, 2001), to incorporate banking and borrowing (Cronshaw and Brown-Kruse, 1996; Rubin, 1996) and to address uncertainty (Schennach, 2000; Zhao, 2003; Newell et al., 2005).

Kosobud et al. (2005) introduced financial tools to the analysis of SO₂ permits in the US Acid Rain program. Other contributions that approach permit markets from a financial perspective rather than the equalization of marginal abatement costs and permit price include Benz and Trueck (2009) and Fehr and Hinz (2006). Seifert et al (2008) explicitly mention the option value of a permit when compared with the alternative of irreversible investment in emissions abatement. Chesney and Taschini (2008) go one step further and define EU ETS allowances to be financial (as opposed to “real”) options. All of these approaches start with the definition of underlying pollution processes, and then derive a market-clearing permit price by way of simulation. In contrast, I take the allowance price series in the EU ETS as given and test whether it is consistent with an options model.

2.2 The European Union Emissions Trading Scheme (EU ETS)

The EU ETS is the world’s largest emissions permit market to date and covers the EU’s carbon dioxide (CO₂) emissions from six industrial sectors (Figure 1), among which power & heat¹ is dominant with about 70% of total emissions. The market is organized into distinct multiyear periods called “phases” that are subject to different rules and emission caps. The first phase of spanned the years 2005-2007 and was considered a pilot run for phase II, which coincides with the Kyoto compliance period of 2008-2012. First-phase allowances (one-time rights to emit one ton of CO₂, denoted as EUA) could not be banked into the second phase and lost their value if unused for compliance.² About 11,000 individual installations received a total of 2.1 billion emission allowances annually, mostly at no cost. For a more detailed discussion of the market setup, see Kruger and Pizer (2004) and the PEW White Paper (2005).

¹ Installations identified by activity code 1 in the Community Independent Transaction Log (CITL). Besides power & heat producers that sell their output on the market this code also includes numerous installations involved in the production of on-site power and heat (“autoproducers”). They account for about 7 % of emissions within activity code 1 (IEA data).

² Banking is allowed between the second and later phases.

Firms can trade allowances freely within the EU. Trades may occur bilaterally, through brokers (over-the-counter or OTC trades) or on one of six exchanges. By April 31, firms have to surrender permits corresponding to their emissions in the previous calendar year. For every ton of emitted CO₂ for which they don't surrender an allowance, they have to pay a penalty (€40 in Phase I and €100 in Phase II) as well as surrender the missing allowance in the following year. Because firms receive annual allowances in March they can effectively bank and borrow across time within a market phase.

(Figure 1 about here)

Figure 2 shows allowance price realizations. The price increased from around €7 in January 2005 to above €30 in April 2006, before crashing to below €10 within three days. It then rose again and stabilized above €15 for about four months before decreasing to practically zero by mid 2007. The April price crash was triggered by the first round of emissions verifications, which revealed that 2005 emissions were 89 MT below the cap. Table 1 shows a summary of the first market phase. The second round of emissions verifications in May 2007 again found an allowance surplus but this had no impact since prices had already decreased to a few cents. Prices in the current phase II are much higher because of a tighter emissions cap and a banking provision into later phases. They are also less volatile, presumably due to improved information about aggregate emissions based on the experiences from phase 1.

(Figure 2 about here)

(Table 1 about here)

2.3 Assumption of stochastic emissions

The assumption that firms cannot perfectly control their emissions and that aggregate emissions are therefore stochastic is central to my analysis. I think that this assumption is appropriate for electricity producers (by far the largest sector covered by the system) during Phase I of the EU ETS for the following reasons:

Because a sufficiently large imbalance of demand and supply crashes the electricity grid leads to large costs from blackouts, many electricity providers are required by law to maintain the grid voltage within a narrow band. The stochastic nature of consumer demand therefore directly translates into stochastic electricity supply.

Given that output is stochastic, emissions are stochastic unless abatement is feasible. There are two possible ways to cut emissions from power production: Building more efficient generators or shifting production from more polluting to less polluting generators within the existing generation portfolio.³ The timely construction of cleaner production technology was largely impossible before the end of the first phase, and borrowing from later phases was not allowed. This leaves essentially only fuel switching as a method of abatement in the power sector (Sijm et al., 2006; Bunn and Fezzi, 2008; Rickels et al., 2007; Mansanet-Bataller et al., 2007; Alberola et al., 2008). However, energy-intensive industries are typically locked into long-term contracts. It is questionable whether firms were able to adjust these contracts in time, and/or whether they were willing to do so, considering the volatility of allowance prices.

Furthermore, even if abatement by fuel switching had been technically feasible, many firms anticipated that their first-phase emissions were going to be used to guide the distribution of second-phase allowances. The European Commission urged member countries not to engage in this sort of allocation “updating”, but to little avail: Most of them based their national allocation plans for Phase II on verified 2005 emissions.⁴ Basing future allocation on current emissions creates a disincentive to abate, because every unit of abatement comes at a cost not only in the current period but also causes a reduction in future free allocation (Boehringer and Lange, 2005). As a result, it is possible that firms stuck to their existing fuel contracts even if they had been able to switch to gas and thus reduce their future free allocation.

³ Note that carbon capture and storage (CSS) was not an option given the existing infrastructure and the time frame.

⁴ One possible reason for this economically highly inefficient choice is the scarcity of information about historic emissions in the EU, which made it politically difficult for the EU member countries to disregard the information gained from the first round of emissions verification.

Lastly, from an empirical perspective, it does not matter whether firms were unable to reduce emissions, or the allowance price simply never reached marginal abatement costs (the cost of fuel switching was far above the allowance price during most of the first phase). In both cases firms would aim to reach compliance exclusively on the permit market while treating emission as a function of stochastic demand.

Options Formula

Let P_t be the closing price for an allowance on day t , with the index $t = (1, 2, \dots, T)$ starting on January 1, 2005 and ending on December 31, 2007. CO₂ emissions on day t are represented by g_t . Let $G_1^t \equiv \sum_{k=1}^t g_k$ denote cumulative realized emissions and $G_t^T \equiv \sum_{k=t+1}^T g_k$ cumulative future emissions until the end of the market. Past emissions G_1^t are observed with certainty but G_t^T is stochastic. Furthermore, let \bar{P} be the penalty for noncompliance and S_0 the total emissions cap over the entire market period imposed by the regulator. Finally, it will be useful to define $S_t = S_0 - G_1^t$ to be the “remaining cap” until the end of the market.

The purchase of an allowance gives the bearer the option to use it for compliance at the end of the period or to sell it. However, if the cap turns out to be not binding, the bearer can retire the allowance. This makes an allowance a financial option, specifically a binary call option, also called a cash-or-nothing option. At time T the payoff from an allowance is:

$$(1) \quad P_T = \begin{cases} 0 & \text{if } S_T > 0 \\ \bar{P} & \text{if } S_T \leq 0 \end{cases}$$

The penalty \bar{P} is the sum €40 and the cost of buying an additional permit for the second phase at a cost of $P_T^{Phase II}$:

$$(2) \quad \bar{P}_T = \text{€}40 + P_T^{\text{Phase II}}$$

At $t < T$, it is not known with certainty whether the cap will be exceeded, provided that it has not been exceeded already. The expected payoff from holding an allowance at time T is

$$(3) \quad \begin{aligned} E_t[P_T | S_t > 0] &= E_t[\bar{P}_T] * \int_{S_t}^{\infty} \zeta_t(G_t^T) dG_t^T \\ E_t[P_T | S_t \leq 0] &= E_t[\bar{P}_T] \end{aligned}$$

where $\zeta_t(G_t^T)$ denotes the probability density function over cumulative future CO₂ emissions and $E_t[\cdot]$ stands for the expectation taken using all information available at time t . Naturally, if the cap is already exceeded at time $t < T$ such that $S_t < 0$, the probability that the allowance price is equal to the penalty is one.

I specify emissions as a linear combination of normally distributed processes (see Section 4), which means that they are normally distributed as well. Options pricing formulae are usually based on log-normally distributed underlying assets, reflecting the idea that total returns are the multiplication of single-period returns. Cumulative emissions, however, are additive rather than multiplicative, and it is therefore appropriate to model them using a normal distribution.⁵ Let μ_t and s_t denote the mean and standard deviation of stochastic cumulative future emissions G_t^T . The variable

$$Q_t \equiv \frac{G_t^T - \mu_t}{s_t} \sim N(0,1)$$

has a standard normal distribution by construction. Let $\varphi(\cdot)$ and $\Phi(\cdot)$ be the probability density function (pdf) and cumulative probability density function (cdf) of the standard normal distribution, respectively. I now convert the integral in (3) into an integral over Q_t :

⁵ In theory, the choice of a normal distribution makes a truncation at zero necessary since negative emissions are not defined. But because CO₂ emissions in the EU are many standard deviations away from zero, the correction implied by the truncation is very small, such that for the remainder of this paper I will neglect the truncation issue.

$$(4) \quad E_t[P_T | S_t > 0] = E_t[\bar{P}_T] * \int_{(S_t - \mu_t)/s_t}^{\infty} \phi(Q_t) dQ_t$$

which is equivalent to

$$(5) \quad E_t[P_T | S_t > 0] = E_t[\bar{P}_T] * \Phi\left(\frac{\mu_t - S_t}{s_t}\right)$$

Arbitrage considerations dictate that the price at time t be equal to the expected price at T, discounted by the risk-free rate of interest r.⁶ This means that the allowance price is a martingale defined by

$$(6) \quad \begin{aligned} P_t | S_t > 0 &= \left(40 e^{-r(T-t)} + P_t^{Phase II}\right) * \Phi\left(\frac{\mu_t - S_t}{s_t}\right) \\ P_t | S_t \leq 0 &= 40 e^{-r(T-t)} + P_t^{Phase II} \end{aligned}$$

where the term in parenthesis is the discounted penalty and the last term is the probability that the cap is exceeded (this last term is unity for the second line). The forward price for phase-2-allowances is already discounted, such that the application of the discount rate r only applies to the cash penalty of €40. Equation (6) is a binary options formula for the allowance price, with the underlying being normally distributed cumulative future emissions.

What remains to be determined in order to evaluate (6) are past emissions and the mean and standard deviation of cumulative future emissions. The latter are not directly observed, but have to be derived from underlying processes and ultimately estimated using market data. This is the subject of the following section.

⁶ Real-world markets are typically not risk-neutral, but option prices based on risk neutrality nevertheless yield the correct (meaning no-arbitrage) solution for traded assets Hull, J. C., 2002. Options, Futures, and Other Derivatives, Ch. 11., Prentice Hall., Upper Saddle River, NJ. . Although risk aversion may be particularly important for the pricing of non-traded assets such as the weather or electricity demand, the price of market risk can never be determined with a sufficient degree of confidence in order to make its inclusion in a pricing formula worthwhile, due to measurement and identification issues (e.g. a greater market fundamental and a higher price of risk have the same effect on the price).

Deriving the mean and standard deviation of future emissions

4.1. CO₂ emissions as a function of exogenous stochastic processes

There exist no data about daily CO₂ emissions, but for the power and heat sector there is something that can serve as a substitute: Daily electricity consumption.

Electricity is special in the sense that demand has to be met with a matching supply at all times in order for the grid not to collapse. I assume that the short-term price elasticity of electricity consumption is zero, such that electricity supply is equal to demand, which in turn is a function of exogenous processes such as the weather and overall economic activity.⁷

Because only generation from fossil fuels contributes to CO₂ emissions, I have to adjust total consumption by the availability of “clean” (i.e. non-CO₂-emitting) sources of electricity, mainly hydroelectric and nuclear power.⁸ Hydroelectric generation depends on rainfall and varies within and between years, but nuclear generation is largely constant due to low marginal but prohibitively high start-up costs.

Let c_t represent overall electricity consumption; c_t^c consumption of conventional fossil-fueled generation; n nuclear power generation (all in Giga-Watt-hours (GWh) per day); and h_t rainfall in the EU in millimeters (mm) per day. Demand for conventional generation is

$$(7) \quad c_t^c = c_t - \eta h_t - n$$

where η is a fixed coefficient translating precipitation into hydroelectric power.⁹ I compute η by dividing the EU’s total hydro generation in 1990-2005 of 4,852,339 GWh by cumulative

⁷ In the long run, consumers will react to higher electricity prices by changing their consumption habits and appliance portfolio, such that electricity demand is also a function of the electricity price. But regardless of the time horizon and the associated energy efficiency of households and industry, exogenous shocks will always drive short-term electricity consumption.

⁸ Although wind generation has increased rapidly during the past few years, it still accounts for a relatively small fraction of total power production.

⁹ Since precipitation can be stored to some extent, either in reservoirs or as snow in the mountains, there is no immediate relationship between precipitation and hydro generation on any given day. On the long run, however, all (net) hydro generation is ultimately due to precipitation, and even though rainfall today may not translate into more generation today, it nevertheless reduces expected conventional generation needed to satisfy consumer electricity demand until the end of the market.

weighted precipitation over the same period of 9,775.28 mm, using installed hydroelectric capacity per country as weights. This results in a conversion factor of $\eta=496.389$ GWh/mm.

In the EU, 12 member countries have nuclear power plants (BE, CZ, DE, ES, FI, FR, HU, NL, SK, SL, SW, UK). Their average total output in the years 2003-2005 was 2,679 GWh per day, which I will use as a measure for n .

The emission intensity (in CO₂/GWh) of the marginal generator varies over the dispatch order (the sequence according to which generators come online, usually based on least cost). The theoretically correct way to express emissions in Europe's power & heat sector is

$$(8) \quad g_t = \int_0^{c_t^c} \Psi_t(y) dy$$

where $\Psi_t(c_t^c)$ is a function transforming conventional thermal power generation into CO₂ emissions. To compute the integral in (8) I would need to know the dispatch order for each day, as well as the marginal emission intensity of all generators involved, information which is not readily available. Instead, I approximate the emission intensity of all generators that are not continuously running by a linear function. This allows me to express (8) as

$$(9) \quad \begin{aligned} g_t &\approx K + \gamma * c_t^c \\ K &\equiv g^{\min} - \gamma * \min(c_t^c) \end{aligned}$$

The parameter γ is the average emission intensity of fossil-fueled electricity generation beyond minimum generation. For the period under consideration (i.e. the first phase of the EU ETS) I treat γ as fixed.

The adjustment parameter K is the difference between CO₂ emissions associated with minimum thermal generation g^{\min} and the (theoretical) emissions if the emission intensity γ

were applicable to inframarginal generation as well.¹⁰ Combining (7) and (9), emissions are a function of a set of parameters and the two stochastic exogenous processes c_t and h_t :

$$(10) \quad g_t = K + \gamma^*(c_t - \eta h_t - n)$$

At time t , the mean of future CO₂ emissions is defined by

$$(11) \quad \begin{aligned} \mu_t = E_t[G_t^T] &= E_t \left[\sum_{k=t+1}^T K + \gamma^*(c_k - \eta h_k - n) \right] \\ &= (T-t)K + \gamma^* \sum_{k=t+1}^T (E_t[c_k] - \eta E_t[h_k] - n) \end{aligned}$$

The variance (for a derivation see Appendix Result 1) is

$$(12) \quad \begin{aligned} s_t^2 = \text{Var}_t[G_t^T] &= \sum_{k=t+1}^T \text{Var}_t[g_k] + 2 \sum_{k=t+1}^T \sum_{u=k+1}^T \text{Cov}_t[g_k, g_u] \\ &= \gamma^2 \sum_{k=t+1}^T (\text{Var}_t[c_k] - 2\eta \text{Cov}_t[c_k, h_k] + \eta^2 \text{Var}_t[h_k]) \\ &\quad + 2\gamma^2 \sum_{k=t+1}^T \sum_{u=k+1}^T (\text{Cov}_t[c_k, c_u] + \eta^2 \text{Cov}_t[h_k, h_u] - \eta \text{Cov}_t[c_k, h_u] - \eta \text{Cov}_t[h_k, c_u]) \end{aligned}$$

Both expressions are functions of the constants η and n , the parameters K and γ , and the mean, variance and covariance of electricity consumption and precipitation, the derivation of which is the subject of the next subsection.

4.2. Properties of the stochastic processes c_t and h_t

For the definition of the stochastic processes of electricity demand and precipitation, I will draw extensively from a paper by Peter Alaton, Boualem Djehiche and David Stillberger (2002). Although their analysis focuses on pricing a weather option over heating-degree days

¹⁰ In principle, the average emission intensity of inframarginal generation could be greater or smaller than the emission intensity of marginal generation. For example, if inframarginal generation consists to a large part of lignite or anthracite coal power plants, then $K > 0$ because these generators have a greater emission intensity than the marginal generators which are predominantly coal and gas generators. On the other hand, if inframarginal generation consists mainly of generators such as combined cycle gas turbines (CCGTs), then $K < 0$. In the EU, $K > 0$ is more likely given the large number of lignite plants in Germany and the new EU member countries from Eastern Europe.

with the underlying process being temperature, it is very similar in principle to both electricity demand and precipitation, as both are exogenously driven stochastic processes that contain deterministic annual fluctuation and long-term trends. The contribution of my paper is not the derivation of the property of such processes, but the application of these methods to CO₂ allowance pricing.

I will model both electricity consumption and precipitation as diffusion processes¹¹ consisting of a deterministic mean and a stochastic part, and which exhibit mean-reversion.¹² For mathematical tractability, I include the stochastic element in the form of a generalized Wiener process. Combining the processes in the index x , they can be described as

$$(13) \quad dx_t = \left[\frac{dx_t^m}{dt} + a_x * (x_t^m - x_t) \right] dt + \sigma_x [i(t)] dW_t^x ; \quad x = c, h$$

This is known as an Ornstein-Uhlenbeck process with a non-zero mean x_t^m and time-varying volatility (for a general treatment see e.g. Bibby and Sorensen (1995)). The term in brackets represents the drift, followed by the diffusion term defined by the standard Wiener process dW_t^x multiplied by the volatility. The first element of the drift term is due to the fact that mean consumption and precipitation change throughout the year. The mean reversion parameters a_x measure the speed at which the processes revert back to their long-term mean.

I constrain the volatility to be constant within each calendar month, but allow it to differ across months. The index i labels the month to which the time index t refers. I will start this index at 1 in January 1976 and finish at 384 in December 2007. Thus,

¹¹ A diffusion process is the solution to a stochastic differential equation. In particular, it is a continuous-time Markov-process with a continuous sample path. This is a realistic description for electricity consumption and precipitation even though the market and weather data are naturally only available for discrete points in time.

¹² Mean reversion is a commonly observed characteristic in many naturally occurring processes, as they generally do not grow without bounds and eventually return to their long-term mean.

$$\begin{aligned}
i(t) &= 1 && \text{if } t \in \text{Jan } 1976 \\
&= 2 && \text{if } t \in \text{Feb } 1976 \\
&\vdots \\
&= 384 && \text{if } t \in \text{Dec } 2007
\end{aligned}$$

Because I assume that the volatility is the same for each calendar month regardless of the year, it follows that $\sigma[i] = \sigma[i + k * 12]$ for any integer k .

I define the long-term mean of electricity consumption and precipitation as

$$\begin{aligned}
(14) \quad c_t^m &= \beta_0^c + \beta_1^c * t + \beta_2^c * \sin[2\pi t / 365 + \omega^c] + D^c * WD_t \\
h_t^m &= \beta_0^h + \beta_1^h * t + \beta_2^h * \sin[2\pi t / 365 + \omega^h]
\end{aligned}$$

The parameters β_0^x and β_1^x ($x = c, h$) describe the level and trend of the two process, respectively, whereas β_2^x describes the amplitudes of the respective sine wave. The phase angles ω^x shift the oscillation of the two processes to their correct position. Lastly, the vector of coefficients D^c (not applicable to rainfall) accounts for differences in electricity consumption across different weekdays, with WD_t being a vector of weekday dummies.

Equation (13) describes two stochastic differential equations. Their solution at time $s \leq t$ is

$$(15) \quad x_t = (x_s - x_s^m) e^{-a_x(t-s)} + x_t^m + \int_s^t e^{-a_x(t-\tau)} \sigma_x[i(\tau)] dW_\tau^x; \quad x = c, h$$

The first term on the RHS is the deviation of actual consumption/precipitation at the present time s from its mean. As time goes on, the impact of this deviation will diminish due to the mean-reversion property of both processes, measured by the exponent. If one of the processes is at its average at time s , or if $t \gg s$, then the first term will drop out, and the expectation at time t simply becomes x_t^m as defined by (14).

The mean and variance of electricity demand and precipitation can be computed as

$$(16) \quad E_s[x_t] = [x_s - x_s^m] e^{-a_x(t-s)} + x_t^m; \quad x = c, h$$

$$\begin{aligned}
(17) \quad \text{Var}_s[x_t] &= E_s \left[\left(x_t - E_s[x_t] \right)^2 \right] = E_s \left[\int_s^t \int_s^t e^{-2a_x*(t-\tau)} \sigma_x^2[i(\tau)] * (dW_\tau^2) \right] \\
&= \int_s^t e^{-2a_x*(t-y)} \sigma_x^2[i(y)] dy; \quad x = c, h
\end{aligned}$$

The second equality follows from the fact that $E[(dW_t^x)^2] = dt$. If the volatility does not change between s and time t , (17) can be solved to

$$(18) \quad \text{Var}_s[x_t] = \frac{\sigma_x^2[i(t)]}{2a_x} \left(1 - e^{2a_x(t-s)} \right); \quad x = c, h; \quad i(s) = i(t)$$

If s and t do not fall within the same month, the expression becomes more complicated. I will denote the first day of each month as $t^{\min}[i(t)] = \min\{t : i(t) = i\}$. In the Appendix (Result 2) I show that for $x = c, h$ and $i(s) \leq i(t)$, the general expression for the variance is

$$(19) \quad \text{Var}_s[x_t] = \frac{1}{2a_x} \left\{ \sum_{k=i(s)}^{i(t)-1} \left(\sigma_x^2[k] - \sigma_x^2[k+1] \right) e^{-2a_x(t-t^{\min}[k+1])} + \sigma_x^2[i(t)] - e^{-2a_x(t-s)} \sigma_x^2[i(s)] \right\}$$

If the volatility is the same for each month, (19) collapses to (18). To calculate the covariance between electricity consumption and rainfall on the same day, note that

$E[dW_t^c dW_t^h] = \rho^{ch} dt$, where $\rho^{ch} \equiv \text{Cov}[c_t, h_t] / \sqrt{\text{Var}[c_t] * \text{Var}[h_t]}$ is the correlation coefficient between the two processes. Thus,

$$\begin{aligned}
\text{Cov}_s[c_t, h_t] &= E_s \left[\left(c_t - E_s[c_t] \right) \left(h_t - E_s[h_t] \right) \right] = E_s \left[\int_s^t \sigma_c[i(\tau)] \sigma_h[i(\tau)] e^{-(a_c+a_h)*(t-\tau)} dW_\tau^c dW_\tau^h \right] \\
&= \int_s^t \rho^{ch} \sigma_c[i(y)] \sigma_h[i(y)] e^{-(a_c+a_h)*(t-y)} dy
\end{aligned}$$

Analogous to the procedure used for the variance, this can be solved to

$$(20) \quad \text{Cov}_s[c_t, h_t] = \frac{\rho^{ch}}{a_c + a_h} \left\{ \sum_{k=i(s)}^{i(t)-1} \left(\sigma_c[k] \sigma_h[k] - \sigma_c[k+1] \sigma_h[k+1] \right) e^{-(a_c+a_h)(t-t^{\min}[k+1])} + \sigma_c[i(t)] \sigma_h[i(t)] - e^{-(a_c+a_h)(t-s)} \sigma_c[i(s)] \sigma_h[i(s)] \right\}$$

Lastly, the covariance between electricity consumption/precipitation on day t and u for $s \leq t \leq u$ is defined by (see Appendix, Result 3):

$$\begin{aligned}
 (21) \quad & Cov_s[x_t, x_u] = e^{-a_x * (u-t)} * Var_s[x_t]; & x = c, h \\
 & Cov_s[c_t, h_u] = e^{-a_h * (u-t)} * Cov_s[c_t, h_t] \\
 & Cov_s[h_t, c_u] = e^{-a_c * (u-t)} * Cov_s[c_t, h_t]
 \end{aligned}$$

I now substitute expressions (16) and (19)-(21) into (11) and (12), which in turn I substitute into (6). The empirical estimation of the parameters $\beta_0^x, \beta_1^x, \beta_2^x, \omega^x, D^c, a_x, \rho^{ch}$ and $\sigma_x(i)$ is the subject of the following section.

Estimation

I evaluate (6) for the period between January 1, 2006 and December 31, 2007, because daily electricity data is not available for most countries before that. To estimate μ_t and s_t I use data through 2005 only, with some exceptions where necessary and as detailed below.

5.1.) Data

Allowance prices: Over-the-counter (OTC) closing allowance prices from Point Carbon.

Electricity consumption: Daily data about electricity consumption is available from the Union for the Coordination of Transmission of Electricity (UCTE)¹³ for continental European countries, including all EU member states except for the Nordic countries,¹⁴ the UK, Ireland, the Baltic States, Malta and Cyprus. Electricity consumption on the third Wednesday of each month¹⁵ is available since 1994 for 9 EU countries, since 1996 for Germany and since 1999 for another 5 countries. Consumption on the weekend following the third Wednesday of each month is available for the year 2000 only. Starting in 2006, electricity consumption is available

¹³ Available at www.ucte.org, last accessed in September 2008.

¹⁴ Sweden, Denmark and Finland. Norway is not part of the EU, and although it is now linked to the EU ETS, this was not the case during the first phase of the market.

¹⁵ Wednesdays are supposed to be the most typical weekdays (as opposed to Mondays and Fridays, which may be slightly different), and the third week is supposed to be the typical week of a month.

on a daily basis for all UCTE countries. To supplement the UCTE data I obtained historic electricity consumption data from the transmission system operators (TSOs) in the UK, Ireland and the Nordic countries.¹⁶ I exclude Malta, Cyprus and the Baltic States due to lack of data. The 20 countries included in the analysis account for 99% of total production in the EU-25. The EU produces nearly all of the electricity it consumes, with net imports/exports accounting for less than 0.1 percent of overall consumption. I therefore exclude imports/exports in my calculations and set EU consumption equal to EU generation. In order to accommodate the variation in type and provenance of the data I will carry out the analyses separately for each group of countries for which the available data is of the same type (e.g. daily vs. monthly) and covers the same time period. The six groups are listed in Table 2. Figures 3a-f show the available pre-2006 electricity consumption data by group.

(Table 2 about here)

(Figure 3 about here)

Precipitation: Data from the European Climate Assessment and Dataset,¹⁷ which contains daily entries for 1,048 monitoring stations located in 42 countries. The period of observation varies from a few years to >150 years, with most series spanning several decades. To model the stochastic process underlying precipitation, I use data covering the years 1976-2005. The conversion of precipitation into hydroelectric power is location-specific. For example, rainfall in the Netherlands or in Denmark is largely irrelevant because these countries have very little installed hydroelectric generation capacity, whereas hydro generation constitutes a large share of total power production in Alpine and Scandinavian countries. I average station entries by

¹⁶ UK: Daily data since 2001 from the National grid, available at <http://www.nationalgrid.com/uk/Electricity/Data/>; Ireland: Daily data since 2002 from Eirgrid, available at <http://www.eirgrid.com>; Denmark: Daily data since 2000 from Energinet, available at <http://www.energinet.dk>; Finland: Daily data since 2004 from Fingrid, available at <http://www.fingrid.fi>; Sweden: Daily data since 2000 from Svenska Kraftnät, available at <http://www.svk.se/web/Page.aspx?id=5794>.

¹⁷ Klein Tank et al. (2007): “Daily Dataset of 20th-Century Surface Air Temperature and Precipitation Series for the European Climate Assessment”, available at eca.knmi.nl, last accessed in September 2008.

country,¹⁸ and then create a weighted European average using installed hydroelectric capacity in 2006 as weights.¹⁹ Installed hydro generation is given in the last column of Table 2.

Weighted precipitation in millimeters (mm) is shown in Figure 4 for a subset of the sample period. Whereas it is difficult to visually discern a pattern in the raw data (Fig. 4a), using moving 7-day-average (Fig. 4b) reveals a clear seasonality.

(Figure 4 about here)

5.2 Parameter estimation for electricity consumption and precipitation

I estimate the parameters $\beta_0^x, \beta_1^x, \beta_2^x, \omega^x, D^x$ and $\sigma_x[i]$ in (14) with a model that features an autoregressive error to account for mean-reversion and multiplicative heteroskedasticity to allow the variance to differ across months:

$$\begin{aligned}
 (22) \quad x_t &= \beta_0^x + \beta_1^x * t + \alpha_1^x * \sin(2\pi t / 365) + \alpha_2^x * \cos(2\pi t / 365) + D^x * WD_t + \varepsilon_t^x \\
 \varepsilon_t^x &= \varphi_x * \varepsilon_{t-1}^x + u_t^x \\
 u_t^x &\sim N(0, \xi_x^2[i(t)]) \\
 \xi_x^2[i(t)] &= \exp\{\lambda_0^x + \lambda_1^x * Jan_t + \dots + \lambda_{11}^x * Nov_t\}; \quad x = c^1, c^2, \dots, c^6, h
 \end{aligned}$$

The index x covers six different electricity consumption series, plus the (weighted) precipitation series, all of which are estimated separately by maximum likelihood. The parameters β_0^x, β_1^x and D^x are the same as in (14) and are estimated directly. The transformation of the sine wave plus the phase angle into a sine and cosine wave is a standard trigonometric relation and serves to linearize the equation. The parameters β_2^x and ω^x can be computed using the estimates of α_1^x and α_2^x :²⁰

¹⁸ For low-lying countries such as Belgium and Luxembourg, I simply take an average of all monitoring stations. However, since hydro generation in the Alps and in Scandinavia is highly location-specific, I take an average of the subset of monitoring stations that are located in or near mountains. A full list of the selected stations is available from the author upon request.

¹⁹ This data comes from UCTE (www.ucte.org) for continental Europe; from Nordpool (www.nordel.org) for Scandinavia; from the Austrian Energy Agency (www.energyagency.at/enercee/) for the Baltic States; from Harrison Harrison, G. P., 2005. Prospects for Hydro in the UK: Between a ROC and a Hard Place? Working paper. University of Edinburgh. for the UK; and from the Electricity Supply Board (ESB, available at http://www.esb.ie/main/about_esb/power_stations_intro.jsp) for Ireland; all accessed in September 2008.

²⁰ See, for example, Beckwith et al. Beckwith, T. G., Marangoni, R. D., Lienhard, J. H. V., 1995. Mechanical Measurements. Addison-Wesley, Reading, Massachusetts, p. 131. The t-statistics and confidence intervals have to be computed using the delta method.

$$(23) \quad \begin{aligned} \beta_2^x &= \sqrt{(\alpha_1^x)^2 + (\alpha_2^x)^2} ; \\ \omega^x &= \arctan[\alpha_2^x / \alpha_1^x] \end{aligned} \quad x = c^1, c^2, \dots, c^6, h$$

I estimate the daily variance $\sigma_x^2[i]$ from the autocorrelation parameters ϕ^x and the variance of the white noise $\xi_x^2[i]$.²¹ For a stationary AR(1) process, the variance is given as

$$(24) \quad \sigma_x^2[i] = E[(x_t - E[x_t])^2] = E[\varepsilon_t^2] = \frac{\xi_x^2[i]}{1 - \phi_x}$$

The mean-reversion parameter a_x measures the speed at which a shock to x_t is felt at later times. From (16), the expectation of future electricity consumption or precipitation is

$$\begin{aligned} E_s[x_t] &= [x_s - x_s^m]e^{-a_x(t-s)} + x_t^m \\ &= \varepsilon_s^x * e^{-a_x(t-s)} + x_t^m \end{aligned} \quad x = c, h$$

This makes it clear that the term $e^{-a_x(t-s)}$ is equivalent to the impulse-response function of the AR(1) process defined by $\chi(t, s) = \phi^{|t-s|}$ (see e.g. Hamilton (1994) p. 53-54), which measures the impact of an exogenous shock occurring in period s on the variable in period t . Equating the two and solving yields

$$(25) \quad a_x = -\ln(\phi_x)$$

All parameter estimates are given in Table 3. I compute the correlation coefficients among the different series ρ^{kl} , $k, l = c^1, c^1, \dots, h$, by using the data from 2006-2007, for which all series have daily entries.²² The results in Table 4 show that electricity consumption across

²¹ Because I cannot estimate an AR(1) parameter with data that only contains entries for every 3rd Wednesday per month, I use the 2006-7 data to estimate this parameter for Series 1-3. Likewise, the estimate of the variance is sensitive to the frequency of measurement Hayashi, T., Yoshida, N., 2005. On covariance estimation of non-synchronously observed diffusion processes. Bernoulli 11, 359-379. and generally improves with greater frequency. I therefore also use the 2006-7 data to estimate the variance and the correlation coefficients (see below). For all other parameters, I use pre-2006 data only. Note that the daily variance and mean reversion parameter for Series 4-6 (for which such a comparison can be made) are not significantly different between pre- and post-2006 data.

²² Hayashi and Yoshida Ibid. developed an unbiased estimator to compute the correlation coefficient between time series of different measuring intervals, but that estimator is not bounded by unity in magnitude, relying on truncation instead. Also, this would only address the problem of differing frequencies within the same time period, but not that of different time periods.

the six different regions is highly correlated, but that precipitation and electricity consumption are not. I will therefore set $Cov_s[c_t^j, h_t] = 0 \quad \forall j$.

(Table 3 about here)

(Table 4 about here)

I derived the expressions for the variance and covariance in (16), (19) and (21) using aggregate electricity consumption. Having six data groups requires the following adjustment:

$$(16') \quad E_s[c_t] = \sum_{j=1}^6 E_s[c_t^j]$$

$$(19') \quad \begin{aligned} Var_s[c_t] &= \sum_{j=1}^6 Var_s[c_t^j] + 2 \sum_{j=1}^6 \sum_{l=j+1}^6 Cov_s[c_t^j, c_t^l] \\ Cov_s[c_t^j, c_t^l] &= \frac{\rho^{jl}}{a_{c^j} + a_{c^l}} \left\{ \sum_{k=i(s)}^{i(t)-1} \left(\sigma_{c^j}[k] \sigma_{c^l}[k] - \sigma_{c^j}[k+1] \sigma_{c^l}[k+1] \right) e^{-(a_{c^j} + a_{c^l})(t - \min\{k+1\})} \right. \\ &\quad \left. + \sigma_{c^j}[i(t)] \sigma_{c^l}[i(t)] - e^{-(a_{c^j} + a_{c^l})(t-s)} \sigma_{c^j}[i(s)] \sigma_{c^l}[i(s)] \right\} \end{aligned}$$

$$(21') \quad Cov_s[c_t, c_u] = \sum_{j=1}^6 e^{-a_{c^j} * (u-t)} * Var_s[c_t^j] + \sum_{j=1}^6 \sum_{l=j+1}^6 \left(e^{-a_{c^j} * (u-t)} + e^{-a_{c^l} * (u-t)} \right) Cov_s[c_t^j, c_t^l]$$

5.3 Parameter estimation in the options formula

I use the obtained results to estimate the options pricing formula. Because emissions were below the total cap at the end of the market as well as for each year individually, I will disregard the second line of equation (6).

The mean and standard deviation of future emissions are a function of free parameters and estimates of the mean, variance and covariance of the processes for electricity consumption and precipitation. Substituting (11) and (12) into (6) and simplifying gives

Using the much higher-frequency data for 2006-2007 is equivalent to assuming that the covariance between electricity consumption in the six different regions and EU-wide weighted precipitation is the same before and after January 1, 2006. For the two groups for which ample data is available (groups 4-6), this assumption appears to hold.

$$P_t = \left(40 e^{-r(T-t)} + P_t^{Phase II} \right) * \Phi \left(\frac{\gamma A_t - (S_0 - TK)}{\gamma B_t} \right)$$

with

$$(26) \quad \begin{aligned} A_t &\equiv \sum_{k=t+1}^T (E_t[c_k] - \eta E_t[h_k] - n) + \sum_{k=1}^t (c_k - \eta h_k - n) \\ B_t &\equiv \left(\sum_{k=t+1}^T Var_t[c_k] + \eta^2 Var_t[h_k] + 2 \sum_{k=t+1}^T \sum_{u=k+1}^T Cov_t[c_k, c_u] + \eta^2 Cov_t[h_k, h_u] \right)^{1/2} \end{aligned}$$

Before proceeding to estimation I have to make two adjustments to (26): First, because my data only covers 2006-2007 emissions from the power and heat sector, I include a parameter V representing 2005 emissions as well as emissions from all other sectors. Firms have expectations about this parameter, but it is evident from the April 2006 price crash that these expectations were updated after the first round of emissions verifications. As a second adjustment I therefore include an adjustment factor V^{EV} multiplied by a dummy variable D_t^{EV} taking the value of zero before, and of one after the first round of emissions verifications.²³

This leads to the following regression specification:

$$(27) \quad \begin{aligned} P_t &= \left(40 e^{-r(T-t)} + P_t^{Phase II} \right) * \Phi \left(\frac{\gamma * A_t - \bar{K} + D_t^{EV} * V^{EV}}{\gamma * B_t} \right) + \varepsilon_t \quad \varepsilon_t \sim N(0, \sigma_\varepsilon^2) \\ \bar{K} &\equiv S_0 - V - \frac{2}{3} * TK \end{aligned}$$

The price crash in April 2006 implies that V^{EV} has to be negative, consistent with firms updating their expectation of the total number of remaining available permits upwards. The combined parameter \bar{K} represents the number of allowances available to firms in the power sector in the years 2006-7, taking into account the difference between inframarginal and marginal emission intensity. The factor $\frac{2}{3}$ that multiplies the correction factor for inframarginal generation stems from the fact that V already contains total emissions for 2005 including this correction. I use an interest rate r of 10% per annum (using 0 and 20% did not

²³ Because there was no sudden price move in either direction after the second round of emissions verifications, I did not introduce a dummy for this event. Likewise, the third round of emissions verifications in May 2008 had no impact on the allowance price.

significantly alter the results) and estimate eq. (27) by nonlinear regression under the assumption that the residuals ε_t are independently and normally distributed with mean zero and variance σ_ε^2 .

In (27), expectations about emissions from other sectors and updating of these expectations after the first round of emission verifications only has an impact on emission levels (through the parameters V and V^{EV}), but not on uncertainty. As an extension I will allow the standard deviation of total future emissions to be different from my estimate B_t for the power sector. Because the uncertainty decreases over time, it makes sense to model this as a multiplicative rather than additive parameter. Lastly, because the first round of emissions verifications may have significantly reduced firms' uncertainty about future aggregate emissions from all sectors, I allow firms' estimate of the emission uncertainty to be updated after the crash:

$$(28) \quad P_t = \left(40 e^{-r(T-t)} + P_t^{Phase II} \right) * \Phi \left(\frac{\gamma * A_t - \bar{K} + D_t^{EV} * V^{EV}}{\gamma * B_t * (\theta + D_t^{EV} * \theta^{EV})} \right)$$

Setting $\theta = 1$ and $\theta^{EV} = 0$ reduces (28) to (27).

The parameters \bar{K} , γ_t and V^{EV} are not individually identified, such that I compute estimates for V^{EV} and γ_t given \bar{K} . For this I choose a sensible range for \bar{K} based on the following calculation:

The total cap S_0 is about 6,300 Mt (million tons) CO₂, or roughly 2,100 Mt per year (Table 1), which I will use for 2005 emissions. Third-party power and CPH producers²⁴ emit about 1,175 Mt per year, leaving 925 Mt for all other industrial emitters in the years 2006-7 (IEA data). If firms' expectations approximately reflect these numbers, then

$V \approx 2,100 \text{ Mt} + 2 * 925 = 3,950 \text{ Mt}$. To get a ballpark number for TK I assume that lignite

²⁴ These are generators that sell all their output on the market. I assume that combined power and heat (CPH) producers optimize their power output and treat heat production as a byproduct. This is of course different for pure heat producers, which I exclude from the analysis and the emissions of which are part of the parameter V .

plants run continuously and are never at the margin, whereas marginal generation consists of a mix of gas and hard coal generation. The emission intensity of lignite is about 225 tCO₂/GWh greater than that of the average remaining fossil fuel generation in Europe, and electric output from lignite is about 290,000 GWh per year (IEA data). This means that

$\frac{2}{3}TK \approx 225 \text{ t/GWh} * 290,000 \text{ GWh} * 2 = 130.5 \text{ Mt}$, and that $\bar{K} \approx 2,220 \text{ Mt}$. I use a range of $1,800 \text{ Mt} \leq \bar{K} \leq 2,600 \text{ Mt}$ to account for the uncertainty embedded in this calculation.

The left panel in Table 5 shows the results from estimating (27) for different values of \bar{K} . The amount of “additional” available permits V^{EV} is 81-116 Mt. This range makes sense in the context of a market that was initially viewed to be tight but then revealed to have an allowance surplus of 89 Mt. The range of the emission intensity γ of 607-876 tCO₂/GWh is also plausible, considering that the emission intensity of gas and coal generators is about 420 and 960 t CO₂/GWh, respectively, and that coal generators are in the majority in Europe. All estimates are statistically significant at $p < 0.001$.

As a measure of model fit I employ the Cox-Snell generalized R^2 , defined by $R^2 = 1 - [L(0) / L(\beta)]^{2/n}$, where $L(\beta)$ and $L(0)$ refer to the likelihood of the full model and of a model that contains only a constant and an emission verification dummy, and n is the number of observations. The generalized R^2 has an intuitive interpretation: It measures the percentage of the variation of the dependent variable that is unexplained by the null model (Nagelkerke, 1991). The resulting R^2 therefore implies that model (27) accounts for 81 % of the allowance price variation that is unexplained by a model that only relies on pre- and post-crash intercepts.

The right panel of Table 5 contains the estimates from specification (28) where I allow firms’ estimate of the standard deviation of future generation to differ from my estimate of B_t , and to be updated after the first round of emissions verifications. The parameter estimate of $\phi = 1.17$ implies that the pre-crash uncertainty of future emissions was 17% greater than the uncertainty that I computed using electricity generation alone. The value of $\phi^{EV} = -0.89$

shows that beliefs about uncertainty were adjusted downwards after the first round of emissions verification, leading to an overall uncertainty estimate that is only 28% of B_t . The regression fit as measured by the generalized R squared increases to 95%.

Whereas the range of the implied emission intensity remains largely the same as in the base specification, the “surprise” number of additionally available permits V^{EV} is now far above -89 Mt for all values of \bar{K} , although still negative. This is consistent with a high pre-crash price if firms expected emissions in 2006-7 to be significantly greater than in 2005 such that the cap would be exceeded even with a modest allowance surplus in the first year. An alternative explanation for the modest values of V^{EV} is that the action in model (28) lies in the uncertainty rather than levels: Updating beliefs about the number of available permits upwards has the same effect on price as revising the uncertainty of future emission downwards (assuming that the system is on track for a nonbinding cap): Both imply that the likelihood that the cap turns out to be binding is decreased, which leads to a crash in the allowance price.

(Table 5 about here)

Figure 5 shows the predicted price series computed using the estimates from the two specifications, along with the actual allowance price. Both models do a reasonable job in tracking the allowance price, although the more flexible second specification follows the EUA much closer. Importantly, both models are able to explain a stabilization of the EUA price at a level significantly above zero after the price crash. This is because although the cap was seen to be generous after the first round of emissions verifications, there remained a nonzero probability of higher-than-expected emissions in the future making the cap binding and the penalty of noncompliance apply. In contrast, models that are based on the equality of permit price and marginal abatement costs would predict a price of zero once the emission were revealed to be below the cap for the first year.

(Figure 5 about here)

The much better model fit of the specification that allows for an update in the uncertainty implies that in the context of stochastic emissions, uncertainty is relatively more important than knowledge about levels in determining the price. This is a common feature in options prices, which depend heavily on uncertainty about the underlying.

Conclusions

In this paper I derive an allowance pricing formula based on the assumption that firms were not able to effectively control their emissions. In this case, emissions are stochastic from the point of view of the firms, and the value of an allowance is characterized by a binary options pricing formula defined by the penalty of noncompliance multiplied by the probability of a binding cap, rather than by marginal abatement costs.

The parameter estimates of the options pricing formula are highly significant and make economic sense, confirming the validity of the model. The predicted allowance price series fits the observed prices well, especially when accounting for uncertainty embedded in emissions from other sectors and allowing firms' expectation of uncertainty to be updated after the first round of emissions verifications. The results imply that uncertainty about future emissions are relatively more important in setting the allowance price than knowledge about emission levels.

Importantly, the model is able to explain the price stabilization after the price crash, followed by a long and steady decline towards zero, which is due to a declining (but nonzero) probability that the cap was going to be binding. Models based on the equality of allowance price and marginal abatement costs would only be able to explain such a movement if marginal abatement costs also expressed a steady decline towards the end of the market, which was almost certainly not the case: Once the market was shown to be over-allocated, total abatement as well as marginal abatement would drop to zero. This may be the reason for the overall poor performance of such models in explaining price drivers in the EU ETS.

My results imply that the allowance price during the first phase of the EU ETS was to a large extent driven by firms hedging against the possibility of having to pay a penalty for noncompliance, rather than by marginal abatement costs. In the context of stochastic emissions and no borrowing, a permit market does not make sense as it is based on the ability of firms to equate marginal abatement costs to permit price and thus control their emissions. It is precisely this reason that makes my assumption about stochasticity of future emissions crucial for my paper. However, even if this assumption is not true, my results can be viewed as a benchmark for the extreme case of no control over emissions. If firms are partially able to control their emissions, the price will still exhibit some options features but at the same time incorporate drivers related to marginal abatement costs.

My results imply that the primary goal of the EU ETS and every other permit market, namely to achieve a given emissions cap at least cost, was not reached during the first phase, and that other instruments such as efficiency improvements on the consumer side (possibly funded by an emissions tax) may have been preferable. However, another goal of the first phase that was probably even more important than that of efficiency was to prepare the market for the Kyoto period, and this seems to have been a success as the EUA price exhibits much less volatility in the current market phase. Naturally, on the long run marginal abatement costs will have to drive prices in the EU ETS, and they may already do so in the current phase that allows for borrowing within a 5-year period and indefinite banking.

References

- Alaton, P., Djehiche, B., Stillberger, D., 2002. On modelling and pricing weather derivatives. *Applied Mathematical Finance* 9, 1-20.
- Alberola, E., Chevallier, J., Cheze, B., 2008. Price drivers and structural breaks in European carbon prices 2005-2007. *Energy Policy* 36, 787-797.
- Beckwith, T. G., Marangoni, R. D., Lienhard, J. H. V., 1995. *Mechanical Measurements*. Addison-Wesley, Reading, Massachusetts.
- Benz, E., Trück, S., 2009. Modeling the price dynamics of CO₂ emission allowances. *Energy Economics* 31, 4-15.
- Bibby, M., Sørensen, M., 1995. Martingale estimation functions for discretely observed diffusion processes. *Bernoulli* 1, 17-39.
- Boehringer, C., Lange, A., 2005. On the Design of Optimal Grandfathering Schemes for Emission Allowances. *European Economic Review* 49, 2041-55.
- Bunn, D. W., Fezzi, C., 2008. A vector error correction model of the interactions among gas, electricity and carbon prices: An application to the cases of Germany and United Kingdom. In: Gulli, F. (Ed), *Markets for carbon and power pricing in Europe: Theoretical issues and empirical analyses*. Edward Elgar Publishing, pp. 145-159.
- Chesney, M., Taschini, L., 2008. The Endogenous Price Dynamics of the Emission Allowances: An Application to CO₂ Option Pricing. *Swiss Finance Institute Research Paper Nr. 08-02*, Zurich.
- Cronshaw, M. B., Brown-Kruse, J., 1996. Regulated Firms in Pollution Permit Markets with Banking. *Journal of Regulatory Economics* 9, 179-89.
- Fehr, M., Hinz, J., 2006. A quantitative approach to carbon price risk modeling. Working paper. ETH Zurich.
- Hamilton, J. D., 1994. Time series analysis. 799, 799.
- Harrison, G. P., 2005. Prospects for Hydro in the UK: Between a ROC and a Hard Place? Working paper. University of Edinburgh.
- Hayashi, T., Yoshida, N., 2005. On covariance estimation of non-synchronously observed diffusion processes. *Bernoulli* 11, 359-379.
- Hull, J. C., 2002. *Options, Futures, and Other Derivatives*, Ch. 11., Prentice Hall, Upper Saddle River, NJ.
- Kosobud, R. F., Stokes, H. H., Tallarico, C. D., Scott, B. L., 2005. Valuing Tradable Private Rights to Pollute the Public's Air. *Review of Accounting and Finance* 4, 50-71.
- Kruger, J., Pizer, W., 2004. The EU Emissions Trading Directive: Opportunities and Potential Pitfalls. *Resources for the Future discussion papers*, Washington, D.C. 62.
- Leiby, P., Rubin, J., 2001. Intertemporal Permit Trading for the Control of Greenhouse Gas Emissions. *Environmental and Resource Economics* 19, 229-56.
- Mansanet-Bataller, M., Pardo, A., Valor, E., 2007. CO₂ Prices, Energy and Weather. *The Energy Journal* 28, 73-92.
- Montgomery, W. D., 1972. Markets in Licenses and Efficient Pollution Control Programs. *Journal of Economic Theory* 5, 395-418.
- Nagelkerke, N. J. D., 1991. A note on a general definition of the coefficient of determination. *Biometrika* 78, 691-692.
- Newell, R., Pizer, W., Zhang, J., 2005. Managing Permit Markets to Stabilize Prices. *Environmental and Resource Economics* 31, 133-57.
- Paolella, M. S., Taschini, L., 2006. An Econometric Analysis of Emission Trading Allowances. *Research Paper Series*. Swiss Finance Institute, Zurich 45.
- PEW Center on Global Climate Change, 2005. *The European Emissions Trading Scheme (EU-ETS): Insights and Opportunities*. 20.
- Rickels, W., Dusch, V., Keller, A., Peterson, S., 2007. The determinants of allowance prices in the European Emissions Trading Scheme - Can we expect an efficient allowance market 2008? . *Kiel Institute for the World Economy Working Paper No. 1387*, 28.
- Rubin, J. D., 1996. A Model of Intertemporal Emission Trading, Banking, and Borrowing. *Journal of Environmental Economics and Management* 31, 269-86.

- Schennach, S. M., 2000. The Economics of Pollution Permit Banking in the Context of Title IV of the 1990 Clean Air Act Amendments. *Journal of Environmental Economics and Management* 40, 189-210.
- Seifert, J., Uhrig-Homburg, M., Wagner, M., 2008. Dynamic Behavior of CO₂ Spot Prices. *Journal of Environmental Economics and Management* 56, 180-194.
- Sijm, J., Neuhoff, K., Chen, Y., 2006. CO₂ cost pass through and windfall profits in the power sector. *Climate Policy* 6, 49-72.
- Zhao, J., 2003. Irreversible Abatement Investment under Cost Uncertainties: Tradable Emission Permits and Emissions Charges. *Journal of Public Economics* 87, 2765-89.

Tables and Figures

Table 1: Summary of Phase I of the EU ETS (Mt=million tons)

	2005	2006	2007	Total Phase I
Price (time average)	€18.40	€18.05	€0.72	€12.39
Trading volume ^a	262 Mt	817 Mt	1,364 Mt	2,443 Mt
Allocation	2,099 Mt	2,072 Mt	2,079 Mt	6,250 Mt
Emissions	2,010 Mt	2,031 Mt	2,041 Mt	6,081 Mt
Surplus (volume)	89 Mt	41 Mt	39 Mt	168 Mt
Surplus (%)	4.22 %	1.98 %	1.85 %	2.69 %

a: OTC and exchange trading for Phase I and II, but excluding bilateral trades

Table 2: Data availability (pre-2006) and installed hydroelectric capacity by country

Country per data series	Start of data series		Source ^a	Hydro capacity in 2006 (MW)
	Type	Year		
Series 1				
Austria	3rd Wed.	1994	UCTE	11,811
Belgium	3rd Wed.	1994	UCTE	1,411
France	3rd Wed.	1994	UCTE	25,457
Greece	3rd Wed.	1994	UCTE	3,133
Italy	3rd Wed.	1994	UCTE	21,070
Luxembourg	3rd Wed.	1994	UCTE	1,128
Netherlands	3rd Wed.	1994	UCTE	37
Portugal	3rd Wed.	1994	UCTE	4,948
Spain	3rd Wed.	1994	UCTE	20,714
Series 2				
Germany	3rd Wed.	1996	UCTE	9,100
Series 3				
Czech Republic	3rd Wed.	1999	UCTE	2,175
Hungary	3rd Wed.	1999	UCTE	46
Poland	3rd Wed.	1999	UCTE	2,324
Slovak Republic	3rd Wed.	1999	UCTE	2,429
Slovenia	3rd Wed.	1999	UCTE	873
Series 4				
UK	daily	2002	Country TSO	4,256
Ireland	Daily	2002	Country TSO	512
Series 5				
Sweden	daily	2001	Country TSO	16,180
Denmark	daily	2000	Country TSO	10
Series 6				
Finland	daily	2004	Country TSO	3,044

a: UCTE: Union for the Coordination of transmission of electricity; TSO: Transmission system operator

Table 3: Parameter estimates for diffusion processes

	c1	c2	c3	c4	c5	c6	h
N	168	144	108	1,460	2,190	730	10,950
Const.	1486.06	1248.56	654.25	763.47	569.68	207.54	23.45
z	22.73	36.44	17.47	16.92	25.12	1.95	44.10
Trend	86.98	5.33	4.84	9.06	-2.07	1.20	-0.01
z	32.44	3.92	3.42	5.57	-2.44	0.33	-0.28
Mo	n/a	n/a	n/a	-20.84	-3.51	0.66	n/a
z	n/a	n/a	n/a	-22.31	-5.98	1.60	n/a
Fr	n/a	n/a	n/a	-20.31	-13.98	1.01	n/a
z	n/a	n/a	n/a	-20.31	-22.31	2.31	n/a
Sa	-416.47	-207.72	-72.13	-128.22	-67.15	-15.71	n/a
z	32.44	3.92	3.42	-101.66	-97.18	-28.56	n/a
Su	-750.21	-328.49	-128.43	-157.64	-72.43	-21.45	n/a
z	-13.80	-26.70	-12.23	-133.32	-103.87	-43.08	n/a
XNY	n/a	n/a	n/a	-86.72	-37.25	-11.54	n/a
z	n/a	n/a	n/a	-20.08	-12.89	-4.85	n/a
b2(sine)	375.85	145.36	116.96	134.10	104.99	36.98	3.00
z	18.99	25.19	32.22	35.09	41.12	10.91	7.06
w(phase)	1.33	1.39	1.41	1.23	1.34	1.35	-0.40
z	42.06	49.94	46.71	38.54	47.09	14.26	-2.96
AR(1)*	0.58	0.39	0.59	0.84	0.86	0.91	0.52
z	18.95	11.32	21.52	95.98	87.02	74.92	103.92
a*	0.54	0.94	0.53	0.18	0.15	0.09	0.65
z	10.24	10.68	11.38	17.15	13.20	6.94	68.02
$\sigma(i)^*$							
Jan	499.71	133.17	88.72	65.21	46.42	23.69	17.21
Feb	316.67	94.92	53.10	45.35	42.27	23.84	16.15
Mar	366.30	119.39	64.96	67.47	41.32	21.19	19.15
Apr	453.41	142.96	79.08	79.32	51.52	31.36	14.49
May	400.48	135.97	55.56	92.75	48.67	33.66	16.28
Jun	387.02	132.80	59.94	45.64	46.51	34.84	16.54
Jul	427.82	116.79	55.50	20.17	30.71	12.12	18.07
Aug	305.51	97.78	50.71	69.65	12.10	7.77	20.91
Sep	389.43	122.02	61.45	23.50	18.56	7.65	20.21
Oct	387.15	120.30	64.95	31.66	27.65	11.17	22.63
Nov	432.38	108.56	73.35	39.00	35.50	19.57	21.50
Dec	414.69	163.85	85.40	96.82	55.91	42.84	17.68

*For series 1-3, based on 2006-7 data; all other estimates based on pre-2006 data

Table 4: Correlation coefficients^a among different series

	c1	c2	c3	c4	c5	c6	h
c1	1.000						
c2	0.8814*	1.000					
c3	0.9016*	0.8730*	1.000				
c4	0.4554*	0.2976*	0.4927*	1.000			
c5	0.5170*	0.3897*	0.6032*	0.9231*	1.000		
c6	0.4588*	0.3672*	0.5573*	0.8496*	0.9418*	1.000	
h	-0.067	0.014	-0.036	-0.038	-0.033	-0.020	1.000

*p<0.05; all coefficients based on 2006-7 data

a: The correlation coefficient between series x_t^i and x_t^j and the corresponding p-value are computed as

$$\hat{\rho} = \frac{\sum_{t=1}^{T^{i,j}} (x_t^i - \bar{x}^i)(x_t^j - \bar{x}^j)}{\sqrt{\sum_{t=1}^{T^{i,j}} (x_t^i - \bar{x}^i)^2} \sqrt{\sum_{t=1}^{T^{i,j}} (x_t^j - \bar{x}^j)^2}} ; \quad p = 2 * \text{ttail} \left(T^{i,j} - 2, |\hat{\rho}| \sqrt{T^{i,j} - 2} / \sqrt{1 - \hat{\rho}^2} \right)$$

where $T^{i,j}$ refers to the number of days for which both series have valid entries.

Table 5: Parameter estimates from estimating (27) and (28); N=513

	Model 1			Model 2			
Cox-Snell Rsq	0.8140			0.9450			
\bar{K}	γ	V^{EV}		γ	V^{EV}	θ	θ^{EV}
(Mt)	(tCO ₂ /GWh)	(Mt)		(tCO ₂ /GWh)	(Mt)		$\theta + \theta^{EV}$
1'800	606.6	-80.6		606.0	-7.8	1.17	-0.89
1'900	640.3	-85.0		639.7	-8.3	1.17	-0.89
2'000	673.9	-89.5		673.3	-8.8	1.17	-0.89
2'100	707.6	-94.0		707.0	-9.2	1.17	-0.89
2'200	741.3	-98.5		740.7	-9.7	1.17	-0.89
2'300	775.0	-103.0		774.3	-10.1	1.17	-0.89
2'400	808.7	-107.0		808.0	-10.5	1.17	-0.89
2'500	842.4	-112.0		841.7	-11.0	1.17	-0.89
2'600	876.1	-116.0		875.3	-11.4	1.17	-0.89
t	983.34	-34.83		463.63	-2.00	3.15	-2.40
p	<0.001	<0.001		<0.001	0.046	0.002	0.017

Figure 1: Annual emissions by sector (in %)




Figure 2: Allowance price and trading volume during phase I of the EU ETS




Figure 3: Available electricity consumption data, pre-2006

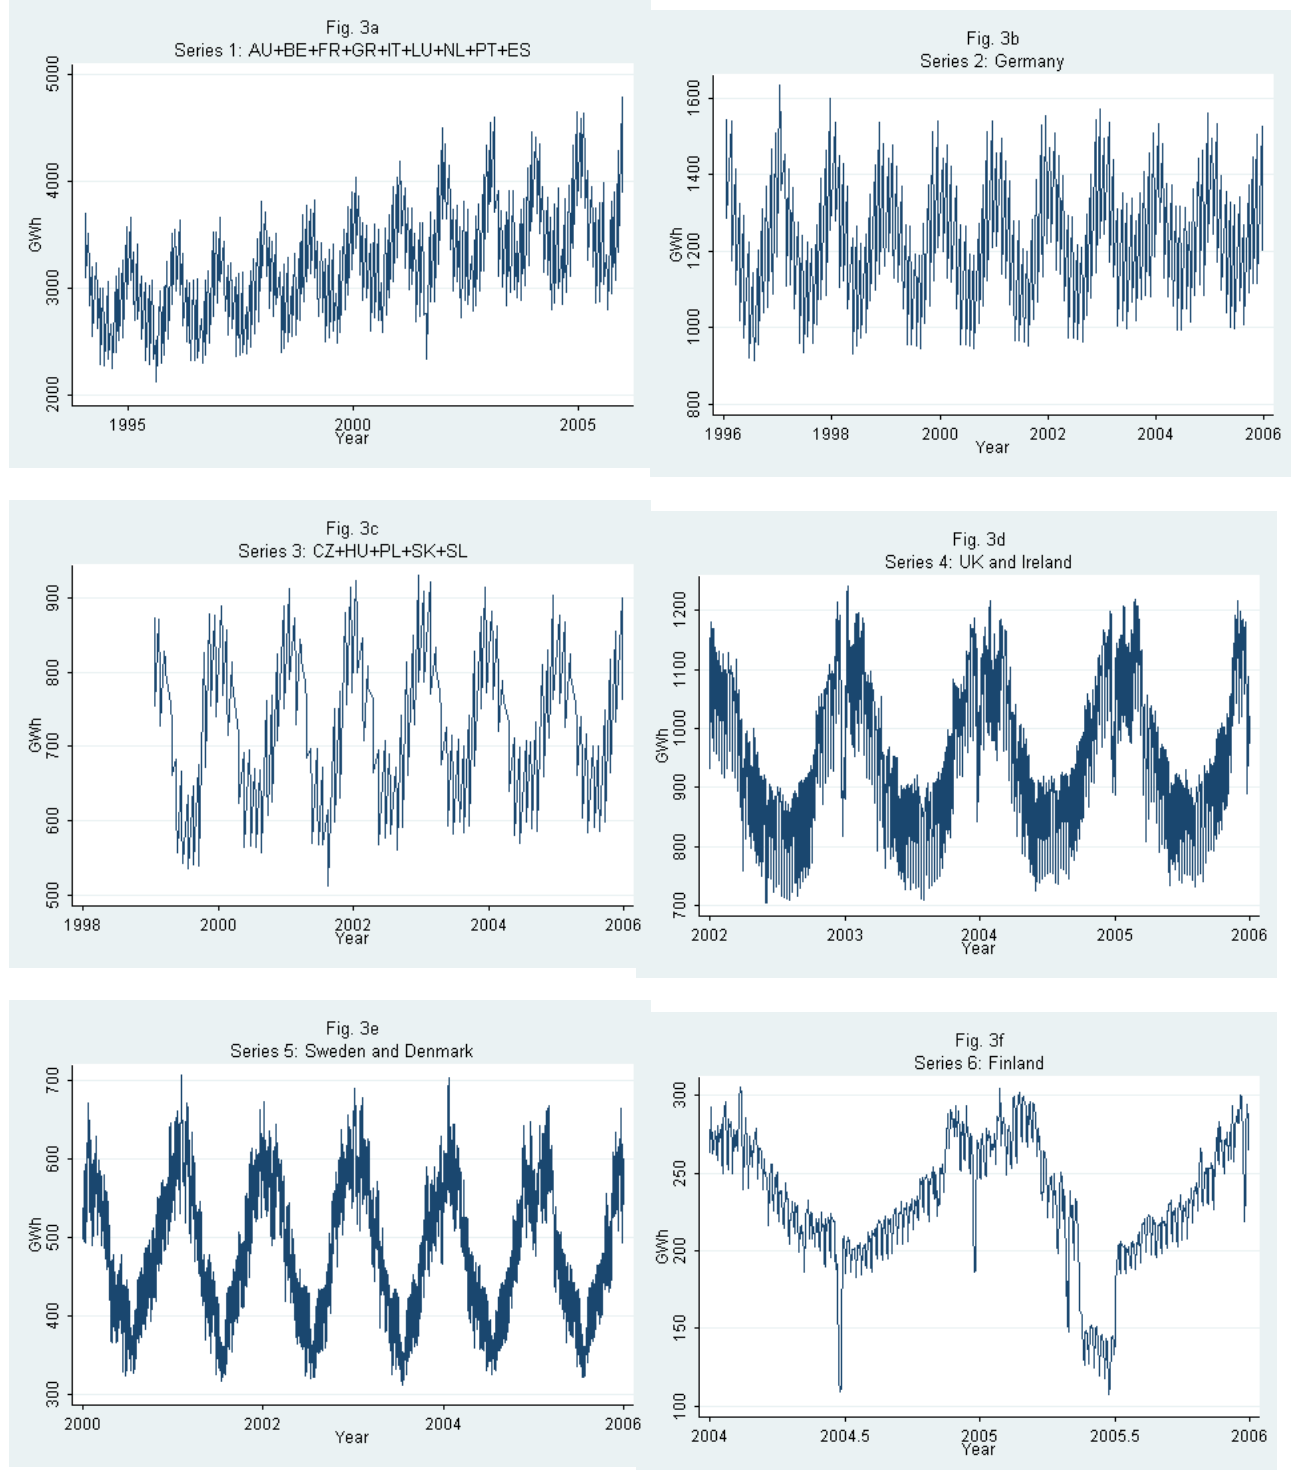


Figure 4: Weighted average precipitation in the EU





Figure 5: EUA price and predictions from estimating (27) and (28)



Appendix

Result 1: Variance of future CO₂ emissions

The variance of G_t^T is defined by

$$(A1) \quad s_t^2 = \text{Var}_s[G_t^T] = \sum_{k=t+1}^T \text{Var}_s[g_k] + 2 \sum_{k=t+1}^T \sum_{u=k+1}^T \text{Cov}_s[g_k, g_u]$$

At time s , the variance g_t and the covariance between g_t and g_u for $s < t < u$ are

$$(A2) \quad \begin{aligned} \text{Var}_s[g_t] &= \text{Var}_s[K + \gamma_1 c_t^c] \\ &= \gamma_1^2 \text{Var}_s[c_t^c] \\ &= \gamma_1^2 (\text{Var}_s[c_t] - 2\eta \text{Cov}_s[c_t, h_t] + \eta^2 \text{Var}_s[h_t]) \end{aligned}$$

$$(A3) \quad \begin{aligned} \text{Cov}_s[g_t, g_u] &= E_s[(g_t - E_s[g_t])(g_u - E_s[g_u])] \\ &= \gamma^2 E_s[\{c_t - E_s[c_t] - \eta(h_t - E_s[h_t])\}\{c_u - E_s[c_u] - \eta(h_u - E_s[h_u])\}] \\ &= \gamma^2 E_s[(c_t - E_s[c_t])(c_u - E_s[c_u])] + \eta^2 \gamma^2 E_s[(h_t - E_s[h_t])(h_u - E_s[h_u])] \\ &\quad - \eta \gamma^2 E_s[(c_t - E_s[c_t])(h_u - E_s[h_u])] - \eta \gamma^2 E_s[(h_t - E_s[h_t])(c_u - E_s[c_u])] \\ &= \gamma^2 (\text{Cov}_s[c_t, c_u] + \eta^2 \text{Cov}_s[h_t, h_u] - \eta \text{Cov}_s[c_t, h_u] - \eta \text{Cov}_s[h_t, c_u]) \end{aligned}$$

Combining (A2) and (A3) establishes the result in equation (12)

$$(12) \quad \begin{aligned} s_t^2 &= \gamma^2 \sum_{k=t+1}^T (\text{Var}_s[c_k] + \eta^2 \text{Var}_s[h_k]) \\ &\quad + 2\gamma^2 \sum_{k=t+1}^T \sum_{u=k+1}^T (\text{Cov}_s[c_k, c_u] + \eta^2 \text{Cov}_s[h_k, h_u] - \eta \text{Cov}_s[c_k, h_u] - \eta \text{Cov}_s[h_k, c_u]) \end{aligned}$$

Result 2: Generalization of the variance for different volatilities

I start by restating the equation (17): The variance of c_t and h_t for $0 \leq s \leq t$ is

$$(17) \quad \text{Var}_s[x_t] = \int_s^t e^{-2a_x(t-\tau)} \sigma_x^2[i(y)] dy \quad x = c, h$$

Suppose that at time s , we're in month 5 and want to calculate the variance of consumption/precipitation in month 8. Using the notation defined in the text that

$t^{\min}[i(t)] = \min\{t : i(t) = i\}$, we have that $s < t^{\min}[6] < t^{\min}[7] < t^{\min}[8] < t < t^{\min}[9]$. I now split up the integral in (17) into four integrals with constant volatility:

$$(A4) \quad Var_s[x_t] = \int_s^{t^{\min}[6]} e^{-2a_x(t-y)} \sigma_x^2[5] dy + \int_{t^{\min}[6]}^{t^{\min}[7]} e^{-2a_x(t-y)} \sigma_x^2[6] dy + \int_{t^{\min}[7]}^{t^{\min}[8]} e^{-2a_x(t-y)} \sigma_x^2[7] dy + \int_{t^{\min}[8]}^t e^{-2a_x(t-y)} \sigma_x^2[8] dy$$

Next, I split the exponents such that they match with the new upper limits of the integrals and move the remainder (a constant) in front:

$$(A5) \quad Var_s[x_t] = e^{-2a_x(t-t^{\min}[6])} \int_s^{t^{\min}[6]} e^{-2a_x(t^{\min}[6]-y)} \sigma_x^2[5] dy + e^{-2a_x(t-t^{\min}[7])} \int_{t^{\min}[6]}^{t^{\min}[7]} e^{-2a_x(t^{\min}[7]-y)} \sigma_x^2[6] dy \\ + e^{-2a_x(t-t^{\min}[8])} \int_{t^{\min}[7]}^{t^{\min}[8]} e^{-2a_x(t^{\min}[8]-y)} \sigma_x^2[7] dy + \int_{t^{\min}[8]}^t e^{-2a_x(t-y)} \sigma_x^2[8] dy$$

Because the volatilities are constant within each integral, each of them can be solved to

$$Var_s[x_t] = e^{-2a_x(t-t^{\min}[6])} * \frac{\sigma_x^2[5]}{2\rho^x} * \left(1 - e^{-2a_x(t^{\min}[6]-s)}\right) + e^{-2a_x(t-t^{\min}[7])} * \frac{\sigma_x^2[6]}{2\rho^x} * \left(1 - e^{-2a_x(t^{\min}[7]-t^{\min}[6])}\right) \\ + e^{-2a_x(t-t^{\min}[8])} * \frac{\sigma_x^2[7]}{2\rho^x} * \left(1 - e^{-2a_x(t^{\min}[8]-t^{\min}[7])}\right) + \frac{\sigma_x^2[8]}{2\rho^x} * \left(1 - e^{-2a_x(t-t^{\min}[7])}\right)$$

Multiplying out and some rearranging gives

$$(A6) \quad Var_s[x_t] = \frac{1}{2a_x} \left\{ \left(\sigma_x^2[5] - \sigma_x^2[6] \right) e^{-2a_x(t-t^{\min}[6])} + \left(\sigma_x^2[6] - \sigma_x^2[7] \right) e^{-2a_x(t-t^{\min}[7])} \right\} \\ + \left(\sigma_x^2[7] - \sigma_x^2[8] \right) e^{-2a_x(t-t^{\min}[8])} + \sigma_x^2[8] e^{-2a_x(t-s)}$$

which can be generalized to

$$(19) \quad Var_s[x_t] = \frac{1}{2a_x} \left\{ \sum_{k=i(s)}^{i(t)-1} \left(\sigma_x^2[k] - \sigma_x^2[k+1] \right) e^{-2a_x(t-t^{\min}[k+1])} + \sigma_x^2[i(t)] - e^{-2a_x(t-s)} \sigma_x^2[i(s)] \right\}$$

Result 3: Covariance of x on two different days

The covariance between x_t and x_u , for $x = c, h$ and $s \leq t \leq u$ is given by

$$\begin{aligned} Cov_s[x_t, x_u] &= E_s \left[(x_t - E_s[x_t])(x_u - E_s[x_u]) \right] \\ &= E_s \left[\int_s^t e^{-\rho^x(t-\tau)} \sigma_x[i(\tau)] dW_\tau * \int_s^u e^{-\rho^x(u-\tau)} \sigma_x[i(\tau)] dW_\tau \right] \end{aligned}$$

I split up the second integral into two parts and pull out the constant term:

$$Cov_s[x_t, x_u] = E_s \left[\int_s^t e^{-\rho^x(t-\tau)} \sigma_x[i(\tau)] dW_\tau * \left(e^{-\rho^x(u-t)} \int_s^t e^{-\rho^x(t-\tau)} \sigma_x[i(\tau)] dW_\tau + \int_t^u e^{-\rho^x(u-\tau)} \sigma_x[i(\tau)] dW_\tau \right) \right]$$

Multiplying out gives

$$\begin{aligned} (A7) \quad Cov_s[x_t, x_u] &= e^{-\rho^x(u-t)} E_s \left[\int_s^t e^{-\rho^x(t-\tau)} \sigma_x[i(\tau)] dW_\tau * \int_s^t e^{-\rho^x(t-\tau)} \sigma_x[i(\tau)] dW_\tau \right] \\ &\quad + E_s \left[\int_s^t e^{-\rho^x(t-\tau)} \sigma_x[i(\tau)] dW_\tau * \int_t^u e^{-\rho^x(u-\tau)} \sigma_x[i(\tau)] dW_\tau \right] \end{aligned}$$

The second term is the expectation of the product of two stochastic processes occurring during non-overlapping time periods. Because a Wiener process is i.i.d., this term drops out. Using the fact that $(dW)^2 = dt$ establishes the result:

$$\begin{aligned} (20) \quad Cov_s[x_t, x_u] &= e^{-\rho^x(u-t)} E_s \left[\int_s^t e^{-\rho^x(t-\tau)} \sigma_x[i(\tau)] dW_\tau * \int_s^t e^{-\rho^x(t-\tau)} \sigma_x[i(\tau)] dW_\tau \right] \\ &= e^{-\rho^x(u-t)} E_s \left[\int_s^t e^{-2\rho^x(t-\tau)} \sigma_x^2[i(\tau)] (dW_\tau)^2 \right] \\ &= e^{-\rho^x(u-t)} * Var_s[x_t] \end{aligned}$$