

Sussex Research

The valuation of clean spread options: linking electricity, emissions and fuels

René Carmona, Michael Coulon, Daniel Schwarz

Publication date

11-12-2012

Licence

This work is made available under the **Copyright not evaluated** licence and should only be used in accordance with that licence. For more information on the specific terms, consult the repository record for this item.

Citation for this work (American Psychological Association 7th edition)

Carmona, R., Coulon, M., & Schwarz, D. (2012). *The valuation of clean spread options: linking electricity, emissions and fuels* (Version 1). University of Sussex. <https://hdl.handle.net/10779/uos.23397002.v1>

Published in

Quantitative Finance

Link to external publisher version

<https://doi.org/10.1080/14697688.2012.750733>

Copyright and reuse:

This work was downloaded from Sussex Research Open (SRO). This document is made available in line with publisher policy and may differ from the published version. Please cite the published version where possible. Copyright and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners unless otherwise stated. For more information on this work, SRO or to report an issue, you can contact the repository administrators at sro@sussex.ac.uk. Discover more of the University's research at <https://sussex.figshare.com/>

THE VALUATION OF CLEAN SPREAD OPTIONS: LINKING ELECTRICITY, EMISSIONS AND FUELS

RENÉ CARMONA, MICHAEL COULON, AND DANIEL SCHWARZ

ABSTRACT. The purpose of the paper is to present a new pricing method for clean spread options, and to illustrate its main features on a set of numerical examples produced by a dedicated computer code. The novelty of the approach is embedded in the use of a structural model as opposed to reduced-form models which fail to capture properly the fundamental dependencies between the economic factors entering the production process.

1. INTRODUCTION

Spread options are most often used in the commodity and energy markets to encapsulate the profitability of a production process by comparing the price of a refined product to the costs of production including, but not limited to, the prices of the inputs to the production process. When the output commodity is electric power, such spread options are called *spark spreads* when the electricity is produced from natural gas, and *dark spreads* when the electricity is produced from coal. Both processes are the sources of GreenHouse Gas (GHG) emissions, in higher quantities for the latter than the former. In this paper we concentrate on the production of electricity and CO₂ emissions and the resulting dependence structure between prices.

Market mechanisms aimed at controlling CO₂ emissions have been implemented throughout the world, and whether they are mandatory or voluntary, cap-and-trade schemes have helped to put a price on carbon in Europe, the US, and around the world. In the academic literature, equilibrium models have been used to show what practitioners have known all along, namely that the price put on CO₂ by the regulation should be included in the costs of production to set the price of electricity (cf. [8]).

Strings of spark spread options (options on the spread between the price of 1MWh of electricity and the cost of the amount of natural gas needed to produce such a MWh) with maturities covering a given period are most frequently used to value the optionality of a gas power plant which can be run when it is profitable to do so (namely when the price of electricity is greater than the cost of producing it), and shut down otherwise. In a nutshell, if an economic agent takes control on day t of a gas power plant for a period $[T_1, T_2]$, then for every day $\tau \in [T_1, T_2]$ of this period, he or she can decide to run the power plant when $P_\tau > h_g S_\tau^g + K$, booking a profit $P_\tau - h_g S_\tau^g - K$ for each unit of power produced, and shut the plant down if $P_\tau \leq h_g S_\tau^g + K$. Here P_τ denotes the price at which one unit (1 MWh) of power can be sold on day τ , S_τ^g the price of one unit of natural gas (typically one MMBtu), h_g the efficiency or heat rate of the plant (i.e. the number of units of natural gas needed to produce one unit of electricity) and K the daily (fixed) costs of operations and maintenance of the plant. Ignoring constraints such as ramp-up rates and start-up costs, this scheduling is also automatically induced when generators bid at the level of their production costs in the day-ahead auction for power. So in this somewhat oversimplified analysis of the optionality of the plant, the value at time t of the control of the plant operation on day τ can be expressed as

Date: May 2012.

Partially supported by NSF - DMS-0739195.

$e^{-r(\tau-t)}\mathbb{E}[(P_\tau - h_g S_\tau^g - K)^+ | \mathcal{F}_t]$, where as usual, the exponent $+$ stands for the positive part, i.e. $x^+ = x$ when $x \geq 0$ and $x^+ = 0$ otherwise, r for the constant interest rate used as discount factor to compute present values of future cash flows, and \mathcal{F}_t denotes the information available on day t when the conditional expectation is actually computed. So the operational control (for example as afforded by a tolling contract) of the plant over the period $[T_1, T_2]$ could be valued on day t as

$$V_t^{PP} = \sum_{\tau=T_1}^{T_2} e^{-r(\tau-t)} \mathbb{E}[(P_\tau - h_g S_\tau^g - K)^+ | \mathcal{F}_t].$$

This rather simplistic way of valuing a power generation asset in the spirit of the theory of real options, had far-reaching implications in the developments of the energy markets, and is the main reason why spread options are of the utmost importance. However, such a valuation procedure is flawed in the presence of emission regulation as the costs of production also have to include the costs specific to the regulation. To be more specific, the day- τ potential profit $(P_\tau - h_g S_\tau^g - K)^+$ of the spark spread has to be modified to $(P_\tau - h_g S_\tau^g - e_g A_\tau - K)^+$ in order to accommodate the cost of the regulation. Here A_τ is the price of one allowance certificate worth one ton of CO₂ equivalent, and e_g is the emission coefficient of the plant, namely the number of tons of CO₂ emitted by the plant during the production of one unit of electricity. Such a spread is often called a *clean spread* to emphasize the fact that the externality is being paid for, and the real option approach to power plant valuation leads to the following *clean price*

$$V_t^{CPP} = \sum_{\tau=T_1}^{T_2} e^{-r(\tau-t)} \mathbb{E}[(P_\tau - h_g S_\tau^g - e_g A_\tau - K)^+ | \mathcal{F}_t]$$

for the control of the plant over the period $[T_1, T_2]$ in the presence of the regulation.

In order to price such cross-commodity derivatives, a joint model is clearly required for fuel prices, electricity prices and carbon allowance prices. Various studies have analyzed the strong links between these price series (cf. [16, 22]). Many reduced-form price models have been proposed for electricity (cf. [4, 17] for a review) with a focus on capturing its stylized features such as seasonality, high volatility, spikes, mean-reversion and fuel price correlation. On the other hand, many authors have argued that these same features are better captured via a structural approach, modelling the dynamics of underlying factors such as demand (load), capacity and fuel prices (early examples include [3, 12, 26, 15]).

Similarly, for carbon emission allowances, exogenously specified processes that model prices directly have been proposed by some (cf. [9]). Others have instead treated the emission process as the exogenously specified underlying factor; in this case the allowance certificate becomes a derivative on cumulative emissions (cf. [27, 13]). However, these models do not take into account the important feedback from the allowance price to the rate at which emissions are produced in the electricity sector — a feature, which is crucial for the justification of any implementation of a cap-and-trade scheme. In a discrete-time framework this feedback mechanism has been addressed, for example in [14] and [8]. In continuous-time the problem has been treated in [6] and [20], whereby the former models the accumulation of emissions as a function of an exogenously specified electricity price process, while the latter uses the bid stack mechanism to infer the emission rate.

The literature on spread options is extensive. In industry, Margrabe's classical spread option formula (cf. [24]) is still widely used, and has been extended by various authors (see [7] for an overview) including to the three commodity case, as required for the pricing of clean spreads (cf. [2]). [10] analyse the pricing of two-asset spread options in a multiscale stochastic volatility model. For electricity markets, pricing formulae for dirty spreads based on structural models have been proposed in [11], in which a closed-form formula is derived in the case of $K = 0$, and in [1], in which semi-closed form formulae are derived for $K \neq 0$ at the expense of a fixed merit order.

The original contributions of the paper are twofold. First, we express the value of clean spread options in a formulation where demand for power and fuel prices are the only factors whose stochastic dynamics are given exogenously, and where the prices of power and emission allowances are derived from a bid-stack based structural model and a forward backward stochastic differential system respectively. The second contribution is the development of a numerical code for the computation of the solution of the pricing problem. First we solve a 4+1 dimensional semilinear partial differential equation to compute the price of an emission allowance, and then we use Monte Carlo techniques to compute the price of the spread option. These computational tools are used to produce the numerical results for the case studies presented in §6 of the paper for the purpose of illustrating the impact of a carbon regulation on the price of spread options. In this final section we first compare the price of spark and dark spread options in two different markets, one with no emission regulation in place and the other governed by an increasingly strict cap-and-trade system. Second, we analyze the impact that different merit order scenarios have on the option prices. Third, we demonstrate the difference between the structural and the reduced-form approach by comparing the option prices produced by our model with those produced by two candidate reduced-form models. Fourth and last, we contrast two competing policy instruments: cap-and-trade, represented by the model we propose, and a fixed carbon tax.

2. THE BID STACK: PRICE SETTING IN ELECTRICITY MARKETS

In order to capture the dependency of electricity price on production costs and fundamental factors in a realistic manner, we use a structural model in the spirit of those reviewed in the recent survey of [5]. The premises of structural models for electricity prices depend upon an explicit construction of the supply curve. Since electricity is sold at its marginal cost, the electricity spot price is given by evaluating the supply function for the appropriate values of factors used to describe the costs of production in the model.

In practice, electricity producers submit day-ahead bids to a central market operator, whose task it is to allocate the production of electricity amongst them. Typically, firms' bids have the form of price-quantity pairs, with each pair comprising the amount of electricity the firm is willing to produce, and the price at which the firm is willing to sell this quantity. Given the large number of generators in most markets, it is common in structural models to approximate the resulting step function of market bids by a continuous increasing curve. Firms' bid levels are determined by their costs of production. An important feature of our model, distinguishing it from most of the commonly-used structural models, is the inclusion, as part of the production costs, of the costs incurred because of the existence of emission regulation.

We assume that, when deciding which firms to call upon to produce electricity, the market operator adheres to the merit order, a rule by which cheaper production units are called upon before more expensive ones. For simplicity, operational and transmission constraints are not considered.

Assumption 1. *The market operator arranges bids according to the **merit order**, in increasing order of production costs.*

The map resulting from ordering market supply in increasing order of electricity costs of production is what is called the bid stack. As it is one of the important building blocks of our model, we define it in a formal way for later convenience.

Definition 1. *The **bid stack** is given by a measurable function*

$$b : [0, \bar{x}] \times \mathbb{R} \times \mathbb{R}^n \ni (x, a, s) \mapsto b(x, a, s) \in \mathbb{R},$$

with the property that for each fixed $(a, s) \in \mathbb{R} \times \mathbb{R}^n$, the function $[0, \bar{x}] \ni x \mapsto b(x, a, s)$ is strictly increasing.

In this definition, $\bar{x} \in \mathbb{R}_{++}$ represents the *market capacity* (measured in MWh) and the variable x the *supply of electricity*. The integer $n \in \mathbb{N} \setminus \{0\}$ gives the number of economic factors (typically the prices, in € say, of the fuels used in the production of electricity), and $s \in \mathbb{R}^n$ the numeric values of these factors. Here and throughout the rest of the paper the *cost of carbon emissions* (measured in € per metric ton of CO₂) is denoted by a . So for a given allowance price, say a , and fuel prices, say s , the market is able to supply x units of electricity at price level $b = b(x, a, s)$ (measured in € per MWh). In other words, $b(x, a, s)$ represents the bid level of the marginal production unit in the event that demand equals x .

The choice of a function b which captures the subtle dependence of the electricity price upon the level of supply and the production costs, is far from trivial, and different approaches have been considered in the literature, as reviewed recently by [5]. In §5.1 we extend the model proposed in [11] to include the cost of carbon as part of the variable costs driving bid levels.

3. RISK-NEUTRAL PRICING OF ALLOWANCE CERTIFICATES

As the inclusion of the cost of emission regulation in the valuation of spread options is the main thrust of the paper, we explain how emission allowances are priced in our model. The model we introduce is close to [20]. However we extend the results found therein to allow the equilibrium bids of generators to be stochastic and driven by fuel prices, a generalization that is vital for our purpose.

We suppose that carbon emissions in the economy are subject to cap-and-trade regulation structured as follows: at the end of the compliance period, each registered firm needs to offset its cumulative emissions with emission allowances or incur a penalty for each excess ton of CO₂ not covered by a redeemed allowance certificate. Initially, firms acquire allowance certificates through free allocation, e.g. through National Allocation Plans (NAP) like in the initial phase of the European Union (EU) Emissions Trading Scheme (ETS), or by purchasing them at auctions like in the Regional Greenhouse Gas Initiative (RGGI) in the North East of the US. Allowances change hands throughout the compliance period. Typically, a firm which thinks that its initial endowment will not suffice to cover its emissions will buy allowances, while firms expecting a surplus will sell them. Adding to these *naturals*, speculators enter the market providing liquidity. Allowances are typically traded in the form of forward contracts and options. In this paper, we denote by A_t the spot price of an allowance certificate maturing at the end of the compliance period. Because their cost of carry is negligible, we treat them as financial products liquidly traded in a market without frictions, and in which long and short positions can be taken.

In a competitive equilibrium, the level of cumulative emissions relative to the cap (i.e. the number of allowance certificates issued by the regulation authority) determines whether — at the end of the compliance period — firms will be subjected to a penalty payment and create a demand for allowance certificates. See [8] for details. For this reason, allowance certificates should be regarded as derivatives on the emissions accumulated throughout the trading period. This type of option written on a non-tradable underlying interest is rather frequent in the energy markets: temperature options are a case in point.

3.1. The Market Emission Rate. As evidenced by the above discussion, the rate at which CO₂ is emitted in the atmosphere as a result of electricity production has to be another important building block of our model. Clearly at any given time, this rate is a function of the amount of electricity produced and because of their impact on the merit order, the variable costs of production, including fuel prices, and notably, the carbon allowance price itself.

Definition 2. *The **market emission rate** is given by a measurable function*

$$\mu_e : [0, \bar{x}] \times \mathbb{R} \times \mathbb{R}^n \ni (x, a, s) \mapsto \mu_e(x, a, s) \in \mathbb{R}_+,$$

on which we will impose necessary technical assumptions later on.

With the definition above, for a given level of electricity supply and for given allowance and fuel prices, $\mu_e = \mu_e(x, a, s)$ represents the rate at which the market emits, measured in tons of CO₂ per hour. Cumulative emissions are then computed by integrating the market emission rate over time. Since any increase in supply can only increase the emission rate, it is of course reasonable from a modelling point of view, to expect μ_e to be increasing as a function of x . Similarly, as the cost of carbon increases, the variable costs (and hence the bids) of pollution intensive generators increase by more than those of environmentally friendlier ones. Dirtier technologies become relatively more expensive and are likely to be scheduled further down in the merit order. As a result cleaner technologies are brought online earlier and we expect μ_e to be decreasing as a function of a .

In §5.2 we propose a specific functional form for μ_e consistent with the bid stack model introduced in §5.1.

3.2. The Pricing Problem. We shall use the following notation. For a fixed time horizon $T \in \mathbb{R}_+$, let $(W_t^0, W_t)_{t \in [0, T]}$ be¹ a $(n + 1)$ -dimensional standard Wiener process on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, $\mathcal{F}^0 := (\mathcal{F}_t^0)$ the filtration generated by W^0 , $\mathcal{F}^W := (\mathcal{F}_t^W)$ the filtration generated by W , and $\mathcal{F} := \mathcal{F}^0 \vee \mathcal{F}^W$ the market filtration. All relationships between random variables are to be understood in the almost surely sense.

Consumers' demand for electricity is given by an \mathcal{F}_t^0 -adapted stochastic process (D_t) taking values in $[0, \bar{x}]$. In response to this demand, producers supply electricity, and we assume that demand and supply are always in equilibrium, so that at any time $t \in [0, T]$ an amount D_t of electricity is supplied. The prices of fuels are observed \mathcal{F}_t^W -adapted stochastic processes (S_t) taking values in \mathbb{R}^n and where $S_t := (S_t^1, \dots, S_t^n)$. As we will see in §3.3, the price of an allowance certificate at time t , say A_t , is now constructed as a \mathcal{F}_t -adapted stochastic process solving a Forward Backward Stochastic Differential Equation (FBSDE). The rate of emission $\mu_e(D_t, A_t, S_t)$ can then be evaluated and the cumulative emissions computed by integrating over time, resulting in a \mathcal{F}_t -adapted process (E_t) .

Since we do not present a calibration of the model to any particular electricity or emissions market, we avoid the difficulties of estimating market prices of risk (see for example [17] and [5] for discussions of some possible ways to approach this thorny issue), and instead choose to specify the dynamics of the processes (D_t) and (S_t) directly under a risk-neutral measure $\mathbb{Q} \sim \mathbb{P}$ chosen by the market for pricing purposes. In practice, various alternative approaches to parameter estimation and calibration could be used to identify a risk-neutral measure which is consistent with liquidly traded products such as power and fuel forward contracts, thereby also estimating the magnitude of risk premia. However, our focus here is on the overall nature of energy price correlations, their structural origins and their important impact on spread option prices, rather than carrying out a study of any one market. For this reason, we also ignore any questions of market incompleteness (ie, the uniqueness of \mathbb{Q}), transaction costs, illiquidity, inelasticity of demand (and also any risk that demand could ever exceed supply, for example causing a black-out), and finally the role of non power sector emissions. While interesting details of the markets, we argue that the inclusion of such effects should not cause any substantial change to the important qualitative conclusions drawn from the model in §6.

3.3. An FBSDE for the Allowance Price. We assume that at time $t = 0$, demand for electricity is known. Thereafter, it evolves according to an Itô diffusion. Specifically, for $t \in [0, T]$, demand

¹To simplify the presentation, from now on, we drop the subscript $t \in [0, T]$, which specifies the time interval on which a stochastic process is defined.

for electricity D_t is the unique strong solution of a stochastic differential equation of the form

$$(1) \quad dD_t = \mu_d(t, D_t)dt + \sigma_d(D_t)d\tilde{W}_t^0, \quad D_0 = d_0 \in (0, \bar{x}),$$

where (\tilde{W}_t^0) is an \mathcal{F}_t^0 -adapted \mathbb{Q} -Brownian motion. The time dependence of the drift allows us to capture the seasonality observed in electricity demand, and the resulting seasonality in prices.

Similarly to demand, the prices of the fuels used in the production process satisfy a system of stochastic differential equations written in vector form as follows:

$$(2) \quad dS_t = \mu_s(S_t)dt + \sigma_s(S_t)d\tilde{W}_t, \quad S_0 = s_0 \in \mathbb{R}^n, \quad t \in [0, T],$$

where (\tilde{W}_t) is an \mathcal{F}_t^W -adapted \mathbb{Q} -Brownian motion. We note that in some cases it may be appropriate to also include time-dependence in the drift or volatility above, in order to capture seasonal patterns in some fuels such as natural gas.

Cumulative emissions are measured from the beginning of the compliance period at time $t = 0$, so that $E_0 = 0$. Subsequently, they are determined by integrating over the market emission rate μ_e introduced in Definition 2. So assuming that the price A_t of an allowance certificate is known, the cumulative emissions process is represented by an absolutely continuous process; i.e. for $t \in [0, T]$,

$$(3) \quad dE_t = \mu_e(D_t, A_t, S_t)dt, \quad E_0 = 0.$$

Note that with this definition the process (E_t) is non-decreasing, which makes intuitive sense considering that it represents a cumulative quantity.

To complete the formulation of the pricing model, it remains to characterize the allowance certificate price process $(A_t)_{t \in [0, T]}$. If our model is to apply to a one compliance period scheme, in a competitive equilibrium, at the end of the compliance period $t = T$, its value is given by a deterministic function of the cumulative emissions:

$$(4) \quad A_T = \phi(E_T),$$

where $\phi : \mathbb{R} \mapsto \mathbb{R}$ is bounded, measurable and non-decreasing. Usually $\phi(\cdot) := \pi \mathbb{I}_{[\Gamma, \infty)}(\cdot)$, where $\pi \in \mathbb{R}_+$ denotes the penalty paid in the event of non-compliance and $\Gamma \in \mathbb{R}_+$ the cap chosen by the regulator as the aggregate allocation of certificates. See [8] for details. Since the discounted allowance price is a martingale under \mathbb{Q} , it is equal to the conditional expectation of its terminal value, i.e.

$$(5) \quad A_t = \exp(-r(T-t)) \mathbb{E}^{\mathbb{Q}}[\phi(E_T) | \mathcal{F}_t], \quad \text{for } t \in [0, T],$$

which implies in particular that the allowance price process (A_t) is bounded. Since the filtration (\mathcal{F}_t) is being generated by the Wiener process, it is a consequence of the Martingale Representation Theorem (cf. [21]) that the allowance price can be represented as an Itô integral with respect to the Brownian motion $(\tilde{W}_t^0, \tilde{W}_t)$. It follows that

$$(6) \quad dA_t = rA_t dt + Z_t^0 d\tilde{W}_t^0 + Z_t \cdot d\tilde{W}_t, \quad \text{for } t \in [0, T],$$

for some \mathcal{F}_t -adapted, square integrable process (Z_t^0, Z_t) .

Combining equations (1), (2), (3), (4) and (6), the pricing problem can be reformulated as the solution of the FBSDE

$$(7) \quad \begin{cases} dD_t = \mu_d(t, D_t)dt + \sigma_d(D_t)d\tilde{W}_t^0, & D_0 = d_0 \in (0, \bar{x}), \\ dS_t = \mu_s(S_t)dt + \sigma_s(S_t)d\tilde{W}_t, & S_0 = s_0 \in \mathbb{R}^n, \\ dE_t = \mu_e(D_t, A_t, S_t)dt, & E_0 = 0, \\ dA_t = rA_t dt + Z_t^0 d\tilde{W}_t^0 + Z_t \cdot d\tilde{W}_t, & A_T = \phi(E_T). \end{cases}$$

Notice that the first two equations are standard stochastic differential equations (in the forward direction of time) which do not depend upon the cumulative emissions and the allowance price.

We will choose their coefficients so that existence and uniqueness of solutions hold. Also, for the sake of convenience, we implicitly assume that the function μ_e , whose first argument was originally restricted to the interval $[0, \bar{x}]$, is defined on the whole $\mathbb{R} \times \mathbb{R} \times \mathbb{R}^n$ by setting $\mu_e(x, a, s) = \mu_e(0, a, s)$ for $x < 0$ and $\mu_e(x, a, s) = \mu_e(\bar{x}, a, s)$ for $x > \bar{x}$. Finally, we make the following assumptions on the coefficients of (7):

Assumption 2. *The functions $\mu_d : [0, T] \times [0, \bar{x}] \hookrightarrow \mathbb{R}$, $\sigma_d : [0, \bar{x}] \hookrightarrow \mathbb{R}$, $\mu_s : \mathbb{R}^n \hookrightarrow \mathbb{R}^n$, $\sigma_s : \mathbb{R}^n \hookrightarrow \mathbb{R}^n \times \mathbb{R}^n$ are such that the first two equations in (7) have a unique strong solution.*

3.4. Existence of a Solution to the Allowance Pricing Problem.

Theorem 1. *If Assumption 3.2 holds, the function μ_e giving the emission rate is Lipschitz with respect to the variable a uniformly in x and s , and $\mu_e(x, 0, s)$ is uniformly bounded in x and s , and the function ϕ giving the terminal condition is bounded, non-decreasing and Lipschitz, then the FBSDE (7) has a unique square integrable solution.*

Proof. Let (D_t) and (S_t) represent the strong solutions of the first two equations of (7) whose existence is guaranteed by Assumption 3.2. These equations being decoupled from the remaining ones, the latter can be treated as a FBSDE with random coefficients and one-dimensional forward and backward components. We claim that existence and uniqueness hold because of Theorem 7.1 of Ma et al. (2011).² Strictly speaking, we cannot apply directly Theorem 7.1 of Ma et al. (2011) because our Wiener process is $(n+1)$ -dimensional. However, a close look at the proof shows that what is really needed is to prove the well-posedness of what the authors call the characteristic BSDE, and the boundedness of its solution and the solutions of the dominating Ordinary Differential Equations (ODE). In the present situation, these equations are rather simple due to the fact that (E_t) has bounded variation, and as a consequence, its volatility vanishes. The two dominating ODEs can be solved explicitly and one can check that the solutions are bounded by inspection. Moreover, the function ϕ giving the terminal condition being uniformly Lipschitz, the characteristic BSDE is one-dimensional, though driven by a multi-dimensional Brownian motion, its terminal condition is bounded, and Kobylanskis comparison results (see the original contribution Kobylanski (2000)) can be used to conclude the proof. \square

The above result is proven for a terminal condition given by a smooth function ϕ . However, as already mentioned earlier, competitive equilibrium arguments for single compliance period models suggest that the function ϕ should be singular. See for example Carmona, R., Fehr, F., Hinz, J., and Porchet, A. (2010). Indeed, in the event of non-compliance, that is, when the cumulative emissions strictly exceed the cap at the end of the compliance period, i.e. when $E_T > \Gamma$, the penalty and the allowance certificate are perfect substitutes; therefore, they ought to have the same monetary value and one should have $A_T = \lambda$ which suggests $\phi(e) = \pi$ whenever $e > \Gamma$. Similarly, in the event of compliance, that is, when the cumulative emissions are strictly below the cap, there will be spare certificates in the market; these certificates will be in zero demand and will therefore expire worthless, so in this case, $A_T = 0$ which suggests $\phi(e) = 0$ whenever $e < \Gamma$. This economic interpretation of the function ϕ giving the terminal condition gives the whole story when the event $\{E_T = \Gamma\}$ has zero probability since we do not have to worry about the definition of $\phi(e)$ when $e = \Gamma$. Hence the importance of knowing if the random variable E_T is continuous (e.g. has a density). Again, see an early discussion of this property in Carmona, R., Fehr, F., Hinz, J., and Porchet, A. (2010), and a systematic analysis in Carmona, Delarue, Espinosa and Touzi (2012), and Carmona and Delarue (2012). We conjecture that a proof in the spirit of the one given in Carmona and Delarue (2012) should work in the setting of this paper as long as μ_e is strictly decreasing in a , providing existence and uniqueness of a solution of the FBSDE when the binary

²We thank Francois Delarue for suggesting this approach.

terminal condition is weakened. Furthermore, Carmona and Delarue also proved that, still under strict monotonicity of μ_e , the aggregate emissions are equal to the cap with positive probability at the end of the compliance period. This shows that the competitive equilibrium argument given earlier is enough to specify a unique emission process (E_t) and a unique price process (A_t) for the allowance, even though the terminal price of an allowance A_T at the end of the compliance period cannot be prescribed *ex ante* on a set of scenarios of positive probability. We suspect that this is also the case in the present situation.

Note added in proof: The conjectured existence and uniqueness of a solution to the FBSDE (7) in our setting was recently proved in Schwarz (2012).

4. VALUING CLEAN SPREAD OPTIONS

In this section we consider the problem of spread option pricing as described in the introduction. Whether the goal is to value a physical asset or to manage the risk associated with financial positions, one needs to compute the price of a European call option on the difference between the price of electricity and the costs of production for a particular power plant. The costs that we take into account are the fixed operation and maintenance costs, the cost of the fuel needed to generate one MWh of electricity and the cost of the ensuing emissions. Letting the \mathcal{F}_t -adapted process (P_t) denote the spot price of electricity, and recasting the informal discussion in the introduction with the notation we chose to allow for several input fuels, a clean spread option with maturity $\tau \in [0, T]$ is characterized by the payoff

$$(P_\tau - h_i S_\tau^i - e_i A_\tau - K)^+,$$

where K represents the value of the fixed operation and maintenance costs, $h_i \in \mathbb{R}_{++}$ and $e_i \in \mathbb{R}_{++}$ denote the specific heat and emission rates of the power plant under consideration, and $S^i \in \{S^1, \dots, S^n\}$ is the price at time τ of the fuel used in the production of electricity. In the special case when S^i is the price of coal (gas) the option is known as a *clean dark (spark) spread* option.

Since we are pricing by expectation, the value V_t^i of the clean spread is given by the conditional expectation under the pricing measure of the discounted payoff; i.e.

$$V_t^i = \exp(-r(\tau - t)) \mathbb{E}^\mathbb{Q} \left[(P_\tau - h_i S_\tau^i - e_i A_\tau - K)^+ | \mathcal{F}_t \right], \quad \text{for } t \in [0, \tau].$$

5. A CONCRETE TWO-FUEL MODEL

We now turn to the special case of two fuels, coal and gas, which differ significantly in their level of emissions per MWh of power generated.

5.1. The Bid Stack. Our bid stack model is a slight variation of the one we proposed in [11]. Here we extend it to include the cost of emissions as part of the variable costs driving firms' bids.

We assume that the coal and gas generators have aggregate capacities \bar{x}_c and \bar{x}_g respectively, so that the market capacity is $\bar{x} = \bar{x}_c + \bar{x}_g$, and their bid levels are given by linear functions of the allowance price and the price of the fuel used for the generation of electricity. We denote these bid functions by b_c and b_g respectively. The coefficients appearing in these linear functions correspond to the marginal emission rate (measured in ton equivalent of CO₂ per MWh) and the heat rate (measured in MMBtu per MWh) of the technology in question. Specifically, for $i \in \{c, g\}$, we assume that

$$(8) \quad b_i(x, a, s) := e_i(x)a + h_i(x)s, \quad \text{for } (x, a, s) \in [0, \bar{x}_i] \times \mathbb{R} \times \mathbb{R},$$

where the *marginal emission rate* e_i and the *heat rate* h_i are given by

$$\begin{aligned} e_i(x) &:= \hat{e}_i \exp(m_i x) \\ h_i(x) &:= \hat{h}_i \exp(m_i x), \quad \text{for } x \in [0, \bar{x}_i]. \end{aligned}$$

Here \hat{e}_i , \hat{h}_i , m_i are strictly positive constants. We allow the marginal emission rate and the heat rate of each technology to vary to reflect differences in efficiencies within the fleet of coal and gas generators. Less efficient plants with higher heat rates have correspondingly higher emission rates. We assume that for each technology the ratio h_i/e_i is fixed, a reasonable approximation which implies that the emissions rate of any coal (gas) plant is simply a fixed multiple of the quantity of coal (gas) burned.

Proposition 1. *With b_c and b_g as above and $I = \{c, g\}$, the market bid stack b is given by*

$$b(x, a, s) = \begin{cases} \left(\hat{e}_i a + \hat{h}_i s_i \right) \exp(m_i x), & \text{if } b_i(x, a, s_i) \leq b_j(0, a, s_j) \text{ for } i, j \in I, i \neq j, \\ \left(\hat{e}_i a + \hat{h}_i s_i \right) \exp(m_i(x - \bar{x}_j)), & \text{if } b_i(x - \bar{x}_j, a, s_i) > b_j(0, a, s_j) \text{ for } i, j \in I, i \neq j, \\ \prod_{i \in I} \left(\hat{e}_i a + \hat{h}_i s_i \right)^{\beta_i} \exp(\gamma x), & \text{otherwise} \end{cases}$$

for $(x, a, s) \in [0, \bar{x}] \times \mathbb{R} \times \mathbb{R}^2$, where $\beta_i = \frac{m_{M \setminus \{i\}}}{m_c + m_g}$ and $\gamma = \frac{m_c m_g}{m_c + m_g}$.

Proof. The proof is a straightforward extension of Corollary 1 in [11]. \square

5.2. The Emission Stack. In order to determine the rate at which the market emits we need to know which generators are supplying electricity at any time. By the merit order assumption the market operator calls upon firms in increasing order of their bid levels. Therefore, given electricity, allowance and fuel prices $(p, a, s) \in \mathbb{R} \times \mathbb{R} \times \mathbb{R}^2$, for $i \in \{c, g\}$, the set of active generators of fuel type i is given by $\{x \in [0, \bar{x}_i] : b_i(x, a, s) \leq p\}$.

Proposition 2. *Assuming that the market bid stack is of the form specified in Proposition 1, the market emission rate μ_e is given by*

$$(9) \quad \mu_e(x, a, s) := \sum_{i \in \{c, g\}} \frac{\hat{e}_i}{m_i} \left(\exp \left(m_i \hat{b}_i^{-1}(b(x, a, s), a, s_i) \right) - 1 \right)$$

for $(x, a, s) \in [0, \bar{x}] \times \mathbb{R} \times \mathbb{R}^2$, where for $i \in \{c, g\}$ we define

$$\hat{b}_i^{-1}(p, a, s_i) := 0 \vee \left(\bar{x}_i \wedge \frac{1}{m_i} \log \left(\frac{p}{\hat{e}_i a + \hat{h}_i s_i} \right) \right),$$

for $(p, a, s) \in \mathbb{R} \times \mathbb{R} \times \mathbb{R}^2$, and as usual, $a \wedge b = \min(a, b)$ and $a \vee b = \max(a, b)$.

Proof. The market emission rate follows from integrating the marginal emission rate e_i for each technology over the corresponding set of active generators and then summing the two. Given the monotonicity of b_i in x and its range $[0, \bar{x}_i]$, the function \hat{b}_i^{-1} describes the quantity of electricity supplied by fuel $i \in \{c, g\}$, and hence the required upper limit of integration. \square

5.3. Specifying the Exogenous Stochastic Factors.

\hat{h}_c	\hat{e}_c	m_c	\bar{x}_c	\hat{h}_g	\hat{e}_g	m_g	\bar{x}_g	\bar{x}
3	0.9	0.00005	12000	7	0.4	0.00003	18000	30000

TABLE 1. Parameters relating to the bid stack and the emission stack.

The Demand Process. We posit that under \mathbb{Q} , the process (D_t) satisfies for $t \in [0, T]$ the stochastic differential equation

$$dD_t = -\eta (D_t - \bar{D}(t)) dt + \sqrt{2\eta\hat{\sigma}D_t(\bar{x} - D_t)}d\tilde{W}_t, \quad D_0 = d_0 \in (0, \bar{x}),$$

where $[0, T] \ni t \mapsto \bar{D}(t) \in (0, \bar{x})$ is a deterministic function giving the level of mean reversion of the demand and $\eta, \hat{\sigma} \in \mathbb{R}_{++}$ are constants. With this definition (D_t) is a Jacobi diffusion process; it has a linear, mean-reverting drift component and degenerates on the boundary. Moreover, subject to $\min(\bar{D}(t), \bar{x} - \bar{D}(t)) \geq \bar{x}\hat{\sigma}$, for $t \in [0, T]$, the process remains within the interval $(0, \bar{x})$ at all times (cf. [18]). To capture the seasonal character of demand, we choose a function $\bar{D}(t)$ of the form:

$$\bar{D}(t) := \varphi_0 + \varphi_1 \sin(2\pi\vartheta t),$$

where the coefficients will be chosen below.

The Fuel Price Processes. We assume that the prices of coal (S_t^c) and gas (S_t^g) follow correlated exponential (or geometric) Ornstein Uhlenbeck processes under the measure \mathbb{Q} ; i.e., for $i \in \{c, g\}$ and $t \in [0, T]$,

$$dS_t^i = -\eta_i \left(\log S_t^i - \bar{s}_i - \frac{\hat{\sigma}_i^2}{2\eta_i} \right) S_t^i dt + \hat{\sigma}_i S_t^i d\tilde{W}_t^i, \quad S_0^i = s_0^i \in \mathbb{R}_{++},$$

where $d\langle W^c, W^g \rangle_t = \rho dt$.

6. NUMERICAL ANALYSIS

We now turn to the detailed analysis of the model we propose. For this purpose we consider a number of case studies in §6.2 to §6.5. To produce the following results we used the numerical schemes explained in Appendix A and Appendix B.

6.1. Choice of Parameters. The tables in this section summarise the parameters used for the numerical analysis of our model that follows below. We refer to the parameters specified in Tables 1 - 5 as the ‘base case’ and indicate whenever we depart from this choice. Note that our parameter choices do not correspond to a particular electricity market, but that all values are within a realistic realm.

Table 1 summarises the parameters specifying the bid curves. We consider a medium sized electricity market served by coal and gas generators and with gas being the dominant technology. For the marginal emission rates, Table 1 implies that $e_c \in [0.9, 1.64]$ and $e_g \in [0.4, 0.69]$ (both measured in tCO₂ per MWh), so that all gas plants are ‘cleaner’ than all coal plants. For the heat rates, we observe that $h_c \in [3, 5.5]$ and $h_g \in [7, 12]$ (both measured in MMBtu per MWh). Using (9) now with $D_t = \bar{x}$, for $0 \leq t \leq T$, and the assumption that there are 8760 production hours in the year, we find, denoting the maximum cumulative emissions by \bar{e} , that $\bar{e} = 2.13 \times 10^8$.

Table 2 contains the parameters for the demand process (D_t) . We model periodicities on an annual and a weekly time scale and the chosen rate of mean reversion assumes that demand reverts to its (time-dependent) mean over the course of one week.

η	φ_0	φ_1	ϑ	$\hat{\sigma}$	d_0
50	21000	3000	1	0.1	21000

TABLE 2. Parameters relating to the demand process.

η_c	\bar{s}_c	$\hat{\sigma}_c$	s_0^c	η_g	\bar{s}_g	$\hat{\sigma}_g$	s_0^g	ρ
1.5	2	0.5	exp(2)	1.5	2	0.5	exp(2)	0.3

TABLE 3. Parameters relating to the fuel price processes.

π	Γ	T	r
100	1.4×10^8	1	0.05

TABLE 4. Parameters relating to the cap-and-trade scheme.

High Eff. Coal		Low Eff. Coal		High Eff. Gas		Low Eff. Gas	
h_c	e_c	h_c	e_c	h_g	e_g	h_g	e_g
3.5	1.05	5.0	1.5	7.5	0.43	11.5	0.66

TABLE 5. Parameters relating to the spread options.

In Table 3 we summarise the parameters that specify the behavior of the prices of coal and gas. Both are chosen to be slowly mean-reverting, at least in comparison to demand. To ease analysis, we assume that all parameters are identical for the two fuels, including mean price levels, both measured in MMBtu.³

Table 4 defines the cap-and-trade scheme that we assume to be in place. The duration of the compliance period T is measured in years and we set the cap at 70% of the upper bound \bar{e} for the cumulative emissions, in order to incentivise a reduction in emissions. This choice of parameters results in A_0 being approximately equal to $\pi/2$, a value for which there is significant initial overlap between gas and coal bids in the stack. Furthermore, the parameters imply a bid stack structure such that at mean levels of coal and gas prices, $A_t = 0$ pushes all coal bids below gas bids, while for $A_t = \pi$ almost all coal bids are above all gas bids.

Finally, in Table 5 we specify the four spread option contracts used in the base case scenario to represent high and low efficiency coal plants, and high and low efficiency gas plants (Note that low efficiency means dirtier and corresponds to high h_i and e_i and vice versa).

We now consider a series of case studies to investigate various features of the model's results in turn. As the model captures many different factors and effects, this allows us to isolate some of the most important implications. In Case Study I, we investigate the impact on coal and gas plants of different efficiencies of creating an increasingly strict carbon emissions market. In Case Study II, we assess the impact on these plants of changes in initial fuel prices. In Case Study III, we compare spread option prices in our model with two simple reduced-form approaches for A_t , which allows us to better understand the role of key model features such as bid stack driven abatement. Finally

³We note that gas and coal prices are typically quoted in different units, and can often differ by a factor of ten or more. However, in our analysis, as we are not fitting to data, coal and gas only play the role of common representative fuel types (and other possibilities include lignite, oil, etc.). Therefore, our parameter choices reflect typical characteristics of energy price behaviour, not specifics. Much more importantly for our analysis, we require that one fuel be significantly 'cleaner' than the other, and that the relative price levels allow for merit order changes driven by the cap-and-trade market.




FIGURE 1. Cap strictness analysis for high efficiency coal (top left), low efficiency coal (top right), high efficiency gas (bottom left) and low efficiency gas (bottom right): spark and dark spread option values plotted against maturity, for varying levels of the cap Γ . Note that the five equally-spaced cap values from 1.8×10^8 to 1.0×10^8 tons of CO_2 imply initial allowance prices of \$5, \$28, \$52, \$80, and \$94.

in Case Study IV, we consider the overall impact of cap-and-trade markets in the electricity sector, by comparing them with a well-known alternative, a fixed carbon tax.

6.2. Case Study I: Impact of the Emission Market. The first effect that we are interested in studying in the model is the impact of the cap-and-trade market on clean spread option prices, for increasingly strict levels of the cap Γ . At one extreme (when the cap is so generous that $A_t \approx 0$, for all $t \in [0, T]$), results correspond to the case of a market without a cap and trade system, while at the other extreme (when the cap is so strict that $A_t \approx \pi \exp(-r(T - t))$, for all $t \in [0, T]$), there is essentially a very high carbon tax which tends to push most coal generators above gas generators in the stack. It is intuitively clear that higher carbon prices typically lead to higher spark spread option prices and lower dark spread option prices, thus favouring gas plants over coal plants, but the relationships can be more involved as we vary between low and high efficiency plants.

In Figure 1, we compare spread option prices corresponding to different efficiency generators (i.e., to different h_i, e_i in the spread payoff) as a function of maturity τ . ‘High’ and ‘low’ efficiency plant indicates values of h_i, e_i chosen to be near the lowest and highest respectively in the stack, as given by Table 5. Within each of the four subplots, the five lines correspond to five different values of the cap Γ , ranging from very lenient to very strict. We immediately observe in Figure 1 the seasonality in spread prices caused by the seasonality in power demand. This is most striking




FIGURE 2. Power Plant Value (sum of spreads over τ) versus s_0^g for low efficiency (left) and high efficiency (right). ‘High A’ corresponds to $\Gamma = 1 \times 10^8$, ‘mid A’ to $\Gamma = 1.4 \times 10^8$ (base case) and ‘low A’ to $\Gamma = 1.8 \times 10^8$, with corresponding values $A_0 = 94, A_0 = 52, A_0 = 5$

for the low efficiency cases (high h_i, e_i), as such plants would rarely be used in shoulder months, particularly in the case of gas. For low efficiency plants, the relationship with cap level (and corresponding initial allowance price) is as one would expect: a stricter cap greatly increases the value of gas plants and greatly decreases the value of the dirtier coal plants. This is also true for high efficiency gas plants, although the price difference (in percentage terms) for different Γ is less, since these are effectively ‘in-the-money’ options, unlike those discussed above. However, the analysis becomes more complicated for high efficiency coal plants, which tend to be chosen to run in most market conditions, irrespective of emission markets. Interestingly, we find that for these options the relationship with Γ (and hence A_0) can be non-monotonic under certain conditions, particularly for high levels of demand, when the price is set near the very top of the stack. In such cases a stricter cap provides extra benefit for the cleaner coal plants via higher power prices (typically set by the dirtier coal plants on the margin) which outweighs the disadvantage of coal plants being replaced by gas plants in the merit order.

6.3. Case Study II: Impact of Fuel Price Changes. Notice that in Table 3, the initial conditions of both gas and coal have been set to be equal to their long term median levels. We now consider the case of the gas price s_0^g being either above or below its long term level, thus inducing a change in the initial merit order. Given the record low prices of under \$2 recently witnessed in the US natural gas market (due primarily to shale gas discoveries), it is natural to ask how such fuel price variations affect our spread option results. Note, however, that since $\eta_c = \eta_g = 1.5$ (implying a typical mean reversion time of 8 months), by the end of the trading period, the simulated fuel price distributions will again be centred near their mean reversion levels. Thus in this case study, we capture a temporary, not permanent, shift in fuel prices.

In Figure 2, we plot the value of coal and gas power plants, as given by the sum of spread options of all maturities $\tau \in [0, T]$. In the first plot, we consider low efficiency (high h_i and e_i) plants, while in the second we consider high efficiencies. The latter are much more likely to operate each day and to generate profits, and are hence much more valuable than the former. However, they also show different relationships with s_0^g , as illustrated for several different cap levels Γ (like in Case Study I above) which correspond to high, low or medium (base case) values of A_0 .⁴ Firstly, for low efficiency plants (left plot), we observe that the gas plant value is typically decreasing in s_0^g , as

⁴In the first plot, the cases ‘coal - low A’ and ‘gas - high A’ would produce values much higher than the other cases, and hence we instead choose ‘coal - mid A’ and ‘gas - mid A’ in order to illustrate the effects on a single plot.

σ_a	ρ_{ac}	ρ_{ag}	σ_e	ρ_{ec}	ρ_{eg}
0.6	-0.2	0.4	0.006	-0.2	0.2

TABLE 6. Parameters for reduced-form comparisons, treating A_t and E_t as GBMs.

we expect, since higher gas prices tend to push the bids from gas above those from coal, meaning there is less chance that the gas plant will be used for electricity generation. Similarly, coal plant values are typically increasing in s_0^g , as more coal plants will be used. Note however, that for some cases, the curves flatten out, as no more merit order changes are possible. This is particularly true for the coal plant when A_0 is very high (and hence once gas drops below a certain point, the coal plant is almost certainly going to remain more expensive to run than all gas plants) and for the gas plant when A_0 is very low (and hence once coal increases above a certain point, the gas plant is almost certainly going to remain more expensive to run than all coal plants).

We now turn our attention to the high efficiency case (right plot), meaning the relatively cheap and clean plants for each technology. As expected, coal benefits from low values of A_0 (ie, a lenient cap) and gas from high values of A_0 (ie, a strict cap). On the other hand, the relationship with s_0^g is now increasing for all cases plotted except that of a gas plant with high A_0 . While it may seem surprising that for low values of A_0 (or medium though not plotted), the gas plant value increases with s_0^g , this is quite intuitive when one considers that the range of bids from gas generators widens as s_0^g increases, implying that the efficient plants can make a larger profit when the inefficient plants set the power price. Indeed, as demand is quite high on average, and gas is 60 percent of the market, it is likely that these efficient gas plants will almost always be ‘in-the-money’ even if coal is lower in the stack. Only in the case that coal is typically above gas and now marginal (i.e. the high A_0 case) is the value of the gas plant decreasing in s_0^g since the plant’s profit margins shrink as gas and coal bids converge.

6.4. Case Study III: Comparison with Reduced-Form. The second analysis we consider is to compare the results of our structural model for the allowance price, with two other simpler models, both of which belong to the class of ‘reduced-form’ models. The first of these treats the allowance price itself as a simple Geometric Brownian Motion (with drift r under \mathbb{Q}), and hence A_T is lognormal at spread maturity, like S_T^c and S_T^g . The second comparison treats the emission process as a Geometric Brownian Motion (GBM), and retains the digital terminal condition $A_T = \pi \mathbb{I}_{\{E_T \geq \Gamma\}}$. As the drift of (E_t) is then simply a constant (chosen to match the initial value A_0 in the full model), there is no feedback from (A_t) to (E_t) , or in other words, no abatement induced by the allowance price. For any time t , A_t is then given in closed-form by a formula resembling the Black-Scholes digital option price. In order to fully specify the two reduced-form models, we need to choose a volatility parameter σ_a or σ_e for each of the GBMs, as well as correlations ρ_{ac}, ρ_{ag} or ρ_{ec}, ρ_{eg} with the Brownian Motions driving the other exogenous factors, coal and gas prices. All of these parameters are chosen to approximately match the levels of volatility and correlation produced by simulations in the full structural model, and are given in Table 6. Finally, note that in all three models that we compare, the power price is given by the same bid stack function as usual, so our aim is to isolate and evaluate the effect of our more sophisticated framework for the allowance price, in comparison to simpler approaches. The cap throughout is $\Gamma = 1.4 \times 10^8$, the base case.

Figure 3 reveals that the difference between the reduced form models and the full structural model is relatively small for high efficiency gas and coal plants which are typically ‘in-the-money’. In contrast a larger gap appears for low efficiency cases, where the reduced form models significantly overprice spread options relative to the stack model. In particular, the case of lognormal emissions produces much higher prices, especially for dark spreads. The intuition is as follows. In the full model, the bid stack structure automatically leads to lower emissions when the allowance price is

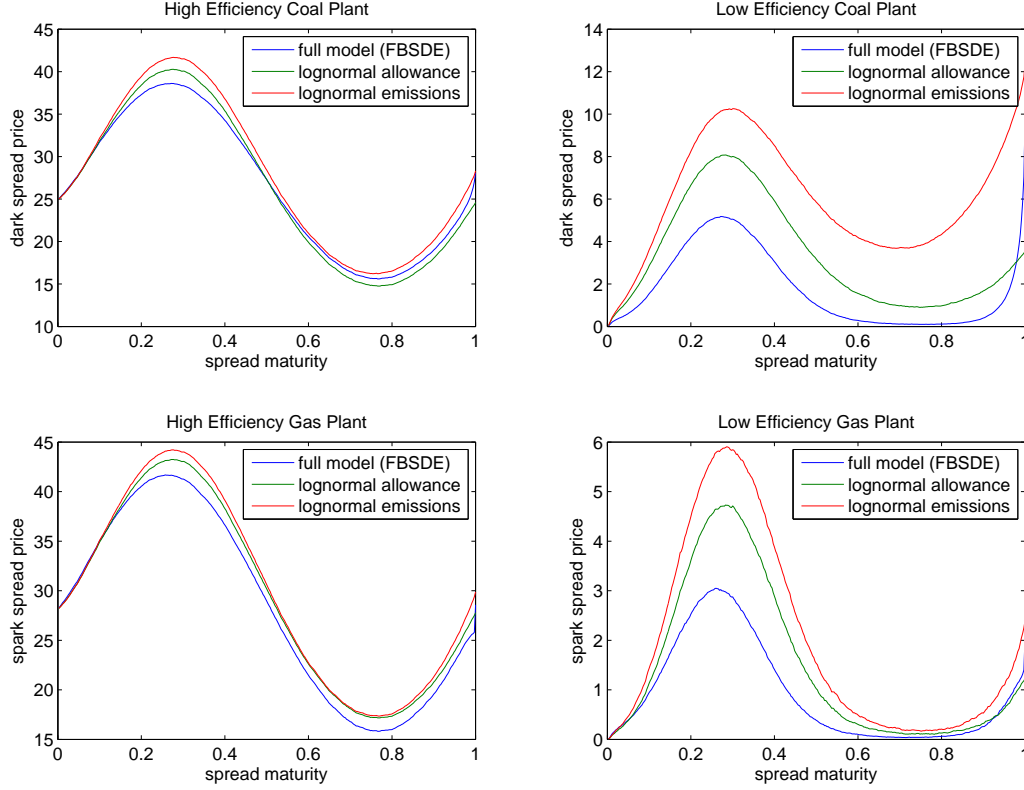


FIGURE 3. Model comparison against reduced-form: Spark and dark spread option values for varying heat rates, emission rates and maturities.

high, and higher emissions when the allowance price is low, producing a mean-reversion-like effect on the cumulative emissions, keeping the process moving roughly towards the cap, with the final outcome (compliance or not) in many simulations only becoming clear very close to maturity. In contrast, if E_t is a GBM, much of the uncertainty is often resolved early in the trading period, with A_t then sticking near zero or π for much of the period. In such cases, there is a much larger benefit for deep OTM options (low efficiency plants), for which the tails of the allowance price distribution provide great value either for coal (when the price is near zero) or for gas (when the price is near the penalty). We observe that in some of the subplots (particularly low efficiency coal), this extra benefit is indeed realized in the full model, but only very near the end of the trading period when the volatility of (A_t) often spikes, and the process either rises or falls sharply. In contrast, for the other reduced-form model with lognormally distributed allowance price, the volatility of the allowance price is constant throughout and (A_t) never moves rapidly towards zero or the penalty. However, the overall link with fuel and power prices is much weaker when simply using correlated Brownian Motions, which serves to widen the spread distribution in most cases relative to the full structural model. This result is somewhat similar to the observation in [11] that a stack model generally produces lower spread option prices than Margrabe's formula for correlated lognormals.

6.5. Case Study IV: Cap-and-Trade vs. Carbon Tax. Finally, we wish to investigate the implications of the model for cap-and-trade systems, as compared with fixed carbon taxes. This question has been much debated by policy makers as well as academics, and can be roughly summarized as fixing quantity versus fixing price. In [8], several different designs for cap-and-trade




FIGURE 4. Cap-and-trade vs. carbon tax: power sector profits versus time for the ‘base case’ (left); total profits over one year for equally-spaced cap values from 1×10^8 to 1.95×10^8 tons of CO_2 (right).

systems are compared to a carbon tax, using criteria such as cost to society and windfall profits to power generators. Here we follow a related approach by analyzing the power sector as a whole, but we build on our previous case studies by using spread option prices as a starting point. Firstly, we observe that the total expected discounted profits of the power sector are equal to the value of all the power plants implied by the bid stack structure, which in turn equals a portfolio of (or integral over) sums of spread option prices with varying h_i and e_i . Therefore, for each simulation over the period $[0, T]$, total profits (total revenues minus total costs) are⁵

$$\begin{aligned}
 \text{Total Profits} &= \sum_{\tau \in [0, T]} \left(P_\tau D_\tau - \int_0^{D_\tau} b(x, A_\tau, S_\tau) dx \right) \\
 &= \sum_{\tau \in [0, T]} \int_0^{\bar{x}} (P_\tau - b(x, A_\tau, S_\tau))^+ dx \\
 &= \sum_{\tau \in [0, T]} \left(\int_0^{\bar{x}_c} (P_\tau - h_c(x) S_\tau^c - e_c(x) A_\tau)^+ dx + \int_0^{\bar{x}_g} (P_\tau - h_g(x) S_\tau^g - e_g(x) A_\tau)^+ dx \right),
 \end{aligned}$$

where the second line follows from the fact that the events $\{P_\tau \geq b(x, A_\tau, S_\tau)\}$ and $\{D_\tau \geq x\}$ are equal.

Hence, instead of picking particular coal and gas plants with efficiencies specified by the parameters in Table 5, we now integrate power plant value over all the efficiencies of plants in the stack, as defined by the parameters in Table 1. In the case of the carbon tax, we simply force $A_t = A_0 \exp(rt)$ for all $t \in [0, T]$, including the exponential function in order to match the mean of the process in the cap-and-trade model. This is equivalent to setting the volatility σ_a equal to zero in the GBM model for the allowance price in Case Study III.

In Figure 4, we first plot the expected total market profits in the base case as a function of time. It is interesting to observe that two important effects occur, pulling the profits in opposite directions, but varying in strength over the trading period. In particular, although the profits must be equal

⁵Note that we do not consider here additional issues such as whether allowances are auctioned or freely allocated to generators. Instead, we assume that allowances are bought on the market by generators as and when they need them.

at time zero, a gap quickly appears in the early part of the trading period, with expected profits to power generators significantly higher under a carbon tax than cap-and-trade. However, as maturity approaches, the gap narrows and the order reverses over the final days, as cap-and-trade generates higher expected profits. We reason as follows: firstly, as $A_0 = 52$ in the base case, the bids of coal and gas begin the period at very similar levels, a state which generally keeps profits low, since the variance of electricity prices is low and the profit margins of both coal and gas generators are quite low. As time progresses and fuel prices move, the coal and gas bids will tend to drift apart in most simulations, for example with gas sometimes moving above coal, say. However, in our structural model for the cap-and-trade scheme, in such a case the higher emissions will induce a higher allowance price, and in turn a feedback effect due to the coupling in (7), which acts to keep coal and gas bids closer together. A similar argument can be made for the case of gas bids tending to move below coal bids but then being counteracted by lower allowance prices. Again we see that the power market structure induces mean reversion on (E_t) , which in this scenario (of an mid-range cap level) corresponds to keeping coal and gas bids close together. On the other hand, under a carbon tax with fixed (or deterministic) allowance price, there is of course no feedback mechanism (no price-sensitive abatement), and bids tend to wander apart. However, as the end of the trading period approaches, in the cap-and-trade system the allowance price eventually gets pulled to either zero or π , which will separate the bids in one way or the other, either leading to very large profits for coal plants (if $A_T = 0$) or for gas plants (if $A_T = \pi$). This is a similar effect to that discussed when comparing with a lognormal allowance price in Case Study III, as neither a carbon tax nor a lognormal allowance price model sees the extra volatility near maturity caused by the terminal condition.

Finally, in the second plot of Figure 4, we consider how these conclusions change if the cap is made stricter or more lenient. Instead of plotting against maturity, we consider the total profits of the power sector over the entire period $[0, T]$. Firstly, we observe that under both forms of emission regulation, power sector profits are lowest if the cap is chosen close to base case, under which the bids from coal and gas generators are more tightly clustered together. Secondly, it is important to notice that the conclusion in the previous discussion that a carbon tax provides more profits to the power sector does not hold for all scenarios of the cap. In particular, for either very high or very low values of the cap, the cap-and-trade scheme provides more profits than a tax. The explanation here is that for the automatic abatement mechanism in the stack to have its largest impact (keeping bids together, and emissions heading towards the cap), there needs to be significant uncertainty at time zero as to whether the cap will be reached. The feedback mechanism of a cap-and-trade system then allows this uncertainty to be prolonged through the period. On the other hand, for an overly strict or overly lenient cap (or similarly for a merit order which does not allow for much abatement), the second effect discussed above dominates over the first. In other words, the terminal condition which guarantees large profits to either coal or gas at maturity begins to take precedence earlier in the trading period, instead of just before maturity as in the base case. Although in practice there are many other details to consider when comparing different forms of emission legislation, our stylized single-period model sheds some light on the differences between cap-and-trade and carbon tax, as well as the clear importance of choosing an appropriate cap level.

7. CONCLUSION

As policy makers debate the future of global carbon emission legislation, the existing cap-and-trade schemes around the world have already significantly impacted the dynamics of electricity prices and the valuation of real assets, such as power plants, particularly under the well-known European Union Emissions Trading Scheme. Together with the recent volatile behaviour of all energy prices (e.g., gas, coal, oil), the introduction of carbon markets has increased the risk of changes in the merit order of fuel types, known to be a crucial factor in the price setting mechanism of electricity

markets. In the US, the recent sharp drop in natural gas prices is already causing changes in the merit order, which would be further magnified by any new emission regulation, such as the upcoming cap-and-trade market in California. Such considerations are vital for describing the complex dependence structure between electricity, its input fuels, and emission allowances, and thus highly relevant for both market participants and policy makers designing emission trading schemes. In this paper, we derived the equilibrium carbon allowance price as the solution of an FBSDE, in which feedback from allowance price on market emission rates is linked to the electricity stack structure. The resulting model specifies simultaneously both electricity and allowance price dynamics as a function of fuel prices, demand and accumulated emissions; in this way, it captures consistently the highly state-dependent correlations between all the energy prices, which would not be achievable in a typical reduced-form approach. We used a PDE representation for the solution of the pricing FBSDE and implemented a finite difference scheme to solve for the price of carbon allowances. Finally we compared our model for allowance prices with other reduced-form approaches and analysed its important implications on price behaviour, spread option pricing and the valuation of physical assets in electricity markets covered by emission regulation. The four case studies illustrated the many important considerations needed to understand the complex joint dynamics of electricity, emissions and fuels, as well as the additional insight that can be provided by our structural approach.

APPENDIX A. NUMERICAL SOLUTION OF THE FBSDE

A.1. Candidate Pricing PDE. In Theorem 1 we addressed the existence and uniqueness of a solution to the FBSDE (7). Given the Markov nature of the equation, we conjecture that there exists a deterministic function $\alpha : [0, T] \times [0, \bar{x}] \times \mathbb{R}_{++} \times [0, \bar{e}] \hookrightarrow [0, \pi]$, such that $A_t = \alpha(t, D_t, E_t, S_t^c, S_t^g)$, and sufficiently smooth to be a classical solution to the semilinear PDE

$$\begin{aligned} (10) \quad & \mathcal{L}\alpha + \mathcal{N}\alpha = 0, & \text{on } U_T \\ (11) \quad & \alpha = \phi(e), & \text{on } \{t = T\} \times U, \end{aligned}$$

where $U := (0, \bar{x}) \times \mathbb{R}_{++} \times \mathbb{R}_{++} \times (0, \bar{e})$ and $U_T := [0, T] \times U$ and the operators \mathcal{L} and \mathcal{N} are defined by

$$\begin{aligned} \mathcal{L} := & \frac{\partial}{\partial t} + \frac{1}{2}\sigma_d(d)^2 \frac{\partial^2}{\partial d^2} + \frac{1}{2}\sigma_c(s_c)^2 \frac{\partial^2}{\partial s_c^2} + \frac{1}{2}\sigma_g(s_g)^2 \frac{\partial^2}{\partial s_g^2} + \rho\sigma_c(s_c)\sigma_g(s_g) \frac{\partial^2}{\partial s_c \partial s_g} \\ & + \mu_d(t, d) \frac{\partial}{\partial d} + \mu_c(s_c) \frac{\partial}{\partial s_c} + \mu_g(s_g) \frac{\partial}{\partial s_g} - r \end{aligned}$$

and $\mathcal{N} := \mu_e(d, \cdot, (s_c, s_g)) \frac{\partial}{\partial e}$. As previously, we specify for our purposes that $\phi(e) = \pi I_{[\Gamma, \infty)}(e)$, for $e \in \mathbb{R}$.

With regards to the problem (10) the question arises at which parts of the boundary we need to specify boundary conditions and, given the original stochastic problem (7), of what form these conditions should be. To answer the former question we consider the Fichera function f at points of the boundary where one or more of the diffusion coefficients disappear (cf. [25]). Defining $n := (n_d, n_c, n_g, n_e)$ to be the inward normal vector to the boundary, Fichera's function for the operator $(\mathcal{N} + \mathcal{L})$ reads

$$\begin{aligned} (12) \quad f(t, d, s_c, s_g, e) := & \left(\mu_d - \frac{1}{2} \frac{\partial}{\partial d} \sigma_d^2 \right) n_d + \left(\mu_c - \frac{1}{2} \frac{\partial}{\partial s_c} \sigma_c^2 - \frac{\partial}{\partial s_c} \rho \sigma_c \sigma_g \right) n_c \\ & + \left(\mu_g - \frac{1}{2} \frac{\partial}{\partial s_g} \sigma_g^2 - \frac{\partial}{\partial s_g} \rho \sigma_c \sigma_g \right) n_g + \mu_e n_e, \quad \text{on } \partial U_T. \end{aligned}$$

At points of the boundary where $f \geq 0$ the direction of information propagation is outward and we do not need to specify any boundary conditions; at points where $f < 0$ information is inward flowing and boundary conditions have to be specified. We evaluate (12) for the choice of coefficients presented in §5.3.

Considering the parts of the boundary corresponding to $d = 0$ and $d = \bar{x}$, we find that $f \geq 0$ if and only if $\min(\bar{D}(t), \bar{x} - \bar{D}(t)) \geq \bar{x}\hat{\sigma}$, which is the same condition prescribed in §5.3 to guarantee that the Jacobi diffusion stays within the interval $(0, \bar{x})$. At points of the boundary corresponding to $e = 0$, we find that $f \geq 0$ always. On the part of the boundary on which $e = \bar{e}$, $f < 0$ except at the point $(d, \cdot, \cdot, e) = (0, \cdot, \cdot, \bar{e})$, where $f = 0$, an ambiguity which could be resolved by smoothing the domain. Similarly, we find that $f \geq 0$ on parts of the boundary where $s_c = 0$ or $s_g = 0$. Therefore, no boundary conditions are necessary except when $e = \bar{e}$, where we prescribe

$$(13) \quad \alpha = \exp(-r(T - t))\pi, \quad \text{on } U_T|_{e=\bar{e}}.$$

In addition we need to specify an asymptotic condition for large values of s_c and s_g . We choose to consider solutions that, for $i \in \{c, g\}$, satisfy

$$(14) \quad \frac{\partial \alpha}{\partial s_i} \sim 0, \quad \text{on } U_T|_{s_i \rightarrow \infty}.$$

A.2. An Implicit - Explicit Finite Difference Scheme. We approximate the domain \bar{U}_T by a finite grid spanning $[0, T] \times [0, \bar{x}] \times [0, \bar{s}_c] \times [0, \bar{s}_g] \times [0, \bar{e}]$. For the discretization we choose mesh widths Δd , Δs_c , Δs_g , Δe and a time step Δt . The discrete mesh points $(t_k, d_m, s_{c_i}, s_{g_j}, e_n)$ are then defined by

$$\begin{aligned} t_k &:= k\Delta t, & d_m &:= m\Delta d, \\ s_{c_i} &:= i\Delta s_c, & s_{g_j} &:= j\Delta s_g, & e_n &:= n\Delta e. \end{aligned}$$

The finite difference scheme we employ produces approximations $\alpha_{m,i,j,n}^k$, which are assumed to converge to the true solution α as the mesh width tends to zero.

Since the partial differential equation (10) is posed backwards in time with a terminal condition, we choose a backward finite difference for the time derivative. In order to achieve better stability properties we make the part of the scheme relating to the linear operator \mathcal{L} implicit; the part relating to the operator \mathcal{N} is made explicit in order to handle the nonlinearity.

In the e -direction we are approximating a conservation law PDE with discontinuous terminal condition. (For an in depth discussion of numerical schemes for this type of equation see [23]) The first derivative in the s -direction, relating to the nonlinear part of the partial differential equation, is discretised against the drift direction using a one-sided upwind difference. Because characteristic information is propagating in the direction of decreasing e , this one-sided difference is also used to calculate the value of the approximation on the part of the boundary corresponding to $e = 0$. At the part of the boundary corresponding to $e = \bar{e}$ we apply the condition (13).

In the d -direction the equation is elliptic everywhere except on the boundary, where it degenerates. Therefore, we expect the convection coefficient to be much larger than the diffusion coefficient near the boundaries. In order to keep the discrete maximum principle we again use a one-sided upwind difference for the first order derivative. Thereby we have to pay attention that due to the mean-reverting nature of (D_t) the direction of information propagation and therefore the upwind direction changes as the sign of μ_d changes. The same upwind difference is also used to calculate the value of the approximation at the boundaries $d = 0$ and $d = \bar{x}$. To discretize the second order derivative we use central differences.

The s_c and s_g -direction are treated similarly to the d -direction. We use one-sided upwind differences for the first order derivatives, thereby taking care of the boundaries corresponding to $s_c = 0$ and

$s_g = 0$. The second order derivatives are discretized using central differences. At the boundary corresponding to $s_c = \bar{s}_c$ and $s_g = \bar{s}_g$ we apply the asymptotic condition (14) as a boundary condition.

With smooth boundary data, on a smooth domain, the scheme described above can be expected to exhibit first order convergence. In our setting, we expect the discontinuous terminal condition to have adverse effects on the convergence rate. We refer the interested reader to [20] for a numerical estimation and a detailed analysis of the convergence rate of the numerical scheme described in this section, as applied to a very similar case.⁶

APPENDIX B. NUMERICAL CALCULATION OF SPREAD PRICES

B.1. Time Discretisation of SDEs. Let $(D_k, S_k^c, S_k^g, E_k, A_k)$ denote the discrete time approximation to the FBSDE solution $(D_t, S_t^c, S_t^g, E_t, A_t)$ on the time grid $0 < \Delta t < 2\Delta t < \dots < n_k \Delta t = \tau$. At each time step we calculate A_k by interpolating the discrete approximation $\alpha_{m,i,j,n}^k$ at (D_k, S_k^c, S_k^g, E_k) , beginning with the initial values $D_0 = d_0, S_0^c = s_0^c, S_0^g = s_0^g, E_0 = 0$. The approximations (D_k, S_k^c, S_k^g, E_k) are obtained using a simple Euler scheme (cf. [19]) for the forward components of (7). The discretized version of (D_t) is forced to be instantaneously reflecting at the boundaries $D_k = 0$ and $D_k = \bar{x}$; similarly, the discretized versions of (S_t^c) and (S_t^g) are made instantaneously reflecting at $S_k^c = 0$ and $S_k^g = 0$.

B.2. Monte Carlo Calculation of Option Prices. Using this discretization we simulate n_{mc} paths and, as usual, for $t \in [0, \tau)$, calculate the mean spark spread price \hat{V}_t^j , given by

$$\hat{V}_t^j := \exp(-r(\tau - t)) \frac{1}{n_{mc}} \sum_{i=1}^{n_{mc}} (b(D_{n_k}^i, S_{n_k}^{c,i}, S_{n_k}^{g,i}, A_{n_k}^i) - h_j S_{n_k}^{g,i} - e_j A_{n_k}^i)^+,$$

where the index i refers to the simulation scenario and $j \in \{c, g\}$. The corresponding standard error $\hat{\sigma}_v$ is obtained by

$$\hat{\sigma}_v := \sqrt{\frac{1}{n_{mc}(n_{mc} - 1)} \sum_{i=1}^{n_{mc}} (V_{n_k}^i - \hat{V}_\tau^j)^2}.$$

REFERENCES

- [1] René Aïd, Luciano Campi, and Nicolas Langrené. A structural risk-neutral model for pricing and hedging power derivatives. *Mathematical Finance*, 2012. to appear.
- [2] E. Alos, A. Eydeland, and P. Laurence. A Kirks and a Bacheliers formula for three-asset spread options. *Energy Risk*, pages 52–57, September 2011.
- [3] Martin Barlow. A diffusion model for electricity prices. *Mathematical Finance*, 12:287–298, 2002.
- [4] F.E. Benth, J.S. Benth, and S. Koekebakker. *Stochastic Modeling of Electricity and Related Markets*, volume 11 of *Advanced Series in Statistical Science & Applied Probability*. World Scientific, 2008.
- [5] R. Carmona and M. Coulon. A survey of commodity markets and structural models for electricity prices. In F. E. Benth, editor, *Financial Engineering for Energy Asset Management and Hedging in Commodity Markets; Proceedings from the special thematic year at the Wolfgang Pauli Institute, Vienna*. 2012.
- [6] R. Carmona, F. Delarue, G.E. Espinoza, and N. Touzi. Singular forward backward stochastic differential equations and emission derivatives. *Annals of Applied Probability*, 2012.
- [7] R. Carmona and V. Durrleman. Pricing and hedging spread options. *SIAM Rev.*, 45(4):627–685, 2003.
- [8] R. Carmona, F. Fehr, J. Hinz, and A. Porchet. Market designs for emissions trading schemes. *SIAM Review*, 52:403–452, 2010.

⁶Although the setting in [20] leads to a $2 + 1$ dimensional model instead of the $4 + 1$ dimensional model we are considering here, we emphasize that structurally the two models are identical. Therefore, we believe that the convergence properties of the two numerical schemes are also similar.

- [9] R. Carmona and J. Hinz. Risk neutral modeling of emission allowance prices and option valuation. Technical report, 2011.
- [10] R. Carmona and Y. Sun. Implied and local correlations from spread options. *Applied Mathematical Finance*, 2012.
- [11] René Carmona, Michael Coulon, and Daniel Schwarz. Electricity price modelling and asset valuation: A multi-fuel structural approach. *Mathematics and Financial Economics*, 2012. To appear.
- [12] Álvaro Cartea and Pablo Villaplana. Spot price modeling and the valuation of electricity forward contracts: the role of demand and capacity. *Journal of Banking and Finance*, 32:2501–2519, 2008.
- [13] M. Chesney and L. Taschini. The endogenous price dynamics of the emission allowances: An application to co2 option pricing. *Applied Mathematical Finance*, 2012.
- [14] Michael Coulon. *Modelling price dynamics through fundamental relationships in electricity and other energy markets*. PhD thesis, University of Oxford, 2009.
- [15] Michael Coulon and Sam Howison. Stochastic behaviour of the electricity bid stack: from fundamental drivers to power prices. *Journal of Energy Markets*, 2:29–69, 2009.
- [16] C. De Jong and S. Schneider. Cointegration between gas and power spot prices. *Journal of Energy Markets*, 2(3):27–46, 2009.
- [17] Alexander Eydeland and Krzysztof Wolyniec. *Energy and Power Risk Management: New Developments in Modeling, Pricing and Hedging*. John Wiley & Sons, 2003.
- [18] Julie L. Forman and Michael Sørensen. The Pearson diffusions: A class of statistically tractable diffusion processes. *Scandinavian Journal of Statistics*, 35:438–465, 2008.
- [19] Paul Glasserman. *Monte Carlo Methods in Financial Engineering*. Springer, 2004.
- [20] Sam Howison and Daniel Schwarz. Risk-neutral pricing of financial instruments in emission markets: A structural approach. *SIAM Journal on Financial Mathematics*, 2012. To appear.
- [21] Ioannis Karatzas and Steven E. Shreve. *Brownian Motion and Stochastic Calculus*. Springer, 1999.
- [22] P. Koenig. Modelling correlation in carbon and energy markets. Working paper, University of Cambridge, Department of Economics, 2011.
- [23] Randall J. LeVeque. *Numerical Methods for Conservation Laws*. Birkhäuser, 1990.
- [24] William Margrabe. The value of an option to exchange one asset for another. *The Journal of Finance*, 33:177–186, 1978.
- [25] O. A. Oleinik and E.V. Radkevich. *Second Order Equations with Nonnegative Characteristic Form*. AMS, 1973.
- [26] Craig Pirrong and Martin Jermakyan. The price of power: The valuation of power and weather derivatives. *Journal of Banking and Finance*, 32:2520–2529, 2008.
- [27] J. Seifert, M. Uhrig-Homburg, and M. W. Wagner. Dynamic behaviour of CO₂ spot prices. *Journal of Environmental Economics and Management*, 56(2):180–194, 2008.

BENDHEIM CENTER FOR FINANCE, DEPT. ORFE, UNIVERSITY OF PRINCETON, PRINCETON NJ 08544, USA

E-mail address: rcarmona@princeton.edu

ORFE, UNIVERSITY OF PRINCETON, PRINCETON NJ 08544, USA

E-mail address: mcoulon@princeton.edu

OXFORD-MAN INSTITUTE, UNIVERSITY OF OXFORD, OXFORD, UK

E-mail address: schwarz@maths.ox.ac.uk