

“Attempts at pricing the regulatory commodity: EU emission credits”

AUTHORS

Peter B. Lerner

ARTICLE INFO

Peter B. Lerner (2010). Attempts at pricing the regulatory commodity: EU emission credits. *Environmental Economics*, 1(1)

RELEASED ON

Friday, 05 November 2010

JOURNAL

"Environmental Economics"

FOUNDER

LLC “Consulting Publishing Company “Business Perspectives”



NUMBER OF REFERENCES

0



NUMBER OF FIGURES

0



NUMBER OF TABLES

0

© The author(s) 2026. This publication is an open access article.

Peter B. Lerner (USA)

Attempts at pricing the regulatory commodity: EU emission credits

Abstract

The market of emission credits and other pollution-related commodities are currently viewed as important instrument in the reduction of environmentally hazardous emissions. European trading in emission credits followed promulgation of the Kyoto treaty by all major EU nations after 1997 and started in earnest in the early 2000s. While the experience with emission credits is widely discussed in its economic, legal and policy implications, pricing of regulation-based commodities remains obscure. Yet, the success of market-based regulations critically depends on the efficiency of prices. This paper proposes a quantitative model that can account for the empirical behavior of EU emission credits in the period of 2005-2007. The paper uses additional information from the December 2008 options prices to formulate the hypothesis that free distribution of permits distorted the behavior of prices on the lower but not on the upper boundary of the price range.

Keywords: ETS, asset pricing models, emission credits, EU permits.

JEL Classification: G12, G13, Q51.

Introduction

Regulatory commodities are nothing new. Hunting permits have been known from time immemorial. English King James I auctioned off the newly-promulgated title of baronet to fill up his treasury (Mayers, 1957) and the sale and purchase of offices was endemic in early modern France (Doyle, 1984). There is a credible view that holy relics were the most liquid form of hoarding wealth during the Middle Age (Duby, 1974; Cambridge, 1987). Recently, the interest in regulatory commodities has increased because many view them as a market-based mechanism to force industrial polluters, to spend money on environmental projects. Commoditized trading is viewed as a favorable alternative to pollution taxes because even within the same industry, enterprises differ greatly in terms of potential cost of environmental cleanups and substantial effect on their core business.

A problem is that environmental cleanup involves significant informational asymmetry between the regulator and the polluter. For the regulator it will be very costly, not only in monetary, but also in political terms (e.g., Axelrod, Vandever and Downie, 2010) to design environmental mitigation projects that can be subsequently tailored to and mandated on a specific firm. An element of trading is intended to compel the polluter spend a maximum “economically reasonable” amount for contributing to the quality of the environment without the need for the regulator to monitor and mandate specific reductions at a given industrial site. The question of “economic reasonableness” refers to the existence of the revealing equilibrium in the game-theoretic sense, which we briefly debate in the next Section.

While some authors maintain that pollution taxes are superior to emission credits (Victor and Cullenward, 2007), there is additional political consideration favoring market-based regulations and cap-and-trade schemes. Namely, pollution-based taxes will always create a feeling with the voting public that as with the cigarette taxes, this is a creative imposition of the government in order to fill its coffers, rather than to address a weighty social problem. Finally, most of the cap-and-trade schemes include tax floors in some form or shape (EU, 2007).

Past experience with regulatory commodities has been mixed. The market for sulfur dioxide emissions has been viewed as a success and a model for subsequent Kyoto-based schemes for emissions and carbon trading (Stavins, 1998). However, the market for installed capacity practiced in several states of the US has been widely regarded as a tax on energy traders in favor of traditional (i.e. generation-owning) utilities, which does not contribute to the efficiency and stability of the energy markets (PMW, 2002; 2005; 2007).

The European Commission initiated the European Climate Change Program in 2000. Its emissions trading scheme (EU ETS) component based on national quotas was launched in 2005. Currently, it is the largest market of environmental permits in the world. We follow first two years of the price history of the ETS trading (2005-2007). Most industrial polluters, such as power utilities throughout the member countries, must receive or purchase carbon allowances to cover their CO₂ output.

The structure of the paper is as follows. In Section 1, I emphasize deficiency of the tax-only approach to pollution control. In Sections 2 and 3 the market for pollution allowances for the period of 2005-2007 (“the first stage”) is described in general terms. In

Section 4, the model is formulated. It is estimated using optimal Kullback-Leibler distance (e.g., Lawler, 2006) in Section 5. In Section 6, I back-test my model on the entire 2005-2008 time series, which include both the first stage of trading and a year of trading after the reallocation of quotas. In Section 7, the additional information from the call and put options on the environmental permits is used to estimate the quality of the model's calibration. Finally, I conclude with the discussion of substantive economic consequences of the paper. In Appendix, I provide the literature review on the pricing of environmental permits, demonstrating the agreement and differences of the presented approach with the approaches existing in the literature.

1. Deficiency of the tax-only approaches to pollution control

The EU policies on climate change do not exclude carbon tax; it is an organic part of an overall emissions reduction mechanism (EU, 2007). The proponents of the carbon tax-only approaches argue that a carbon tax provides better guidelines for the firm and can be enforced more easily (Victor and Cullenward, 2007).

However, there is a consistent microeconomic argument why the tax-only approach is unworkable even if it were feasible, which is doubtful in the current political climate. The tax-only approach provides an incentive for the firm to cheat and underreport pollution. Apprehension and prosecution of cheaters is an expensive and time-consuming process, especially in the view of skeptical state legislatures. Hence, the regulator has an incentive to implicitly include the cheating below the observable threshold in her projected tax rate. This regulatory approach is applied, (e.g., in the allocation of parking spaces). The parking fines are set high enough to deter the violator given an expected probability to escape detection. Yet, this policy works only because there is a simple method to charge a violation to a specific car.

Carbon, however, is elemental. In general, the regulator cannot distinguish the exact contribution of each firm to overall carbon emissions. The firms (two, in the simplest example) can both pay a prevailing tax rate, or choose to cheat. If a firm chooses to pay the tax honestly in Step 0, it must stop production in Step 1 in the high-tax regime set by the regulator, but can maintain its production in a low-tax framework. One can demonstrate that each subgame of this game (firm 1-firm 2, regulator-firm 1 and regulator-firm 2) has a payoff structure similar to a prisoner's dilemma (Hirschleifer, Glazer and

Hirschleifer, 2005) and the game is inherently unstable. The problem is that the regulator in a high-tax regime has a better outcome if some but not all of the firms cheat, because total compliance reduces the tax base. In a low-tax regime she is indifferent between taxes being distributed among all, or just a few firms, but they all benefit from cheating.

The absence of revealing equilibrium, i.e. the equilibrium in which firms benefit from disclosure of the actual, or close to actual, amount of their pollution, devalues the tax-only enforcement. It must necessarily be appended with selective monitoring and high fines and/or prosecution of the cheaters, which is not necessarily simpler or less controversial than the cap-and-trade schemes.

2. Possible approaches to the spot time series

As always in asset pricing, for instance, in the modeling of risky bonds (Duffie and Singleton, 2003), there are two approaches. In the *structural approach*, one builds an economic model of a firm and then values an asset as a specific contingent claim (i.e. a derivative) on the asset's cash flows. In our case this will require the equality of the marginal cost of emission reduction to the marginal cost of remaining emissions to the society (Marshall, 1975; Mas-Colell, Whinston and Green, 1995). We take the other approach, a *reduced-form model*, where the only inputs are the past time series, which are used to calibrate a rather arbitrary model that we hope to reconcile with the economics of the problem.

The conventional assumption is that the structural approach is always better, yet the program it requires is rarely accomplished or is based on a model of economic behavior so simplified as to be almost meaningless. Hence, we choose a reduced-form model, in which we analyze the time series of EU emission credits.

The allocation of permits during the trial period (2005-2007) was widely regarded as long on permits for most of the member nations. In the first phase of the trading scheme, the firms were not allowed to roll forward their permits. If supply consistently outstrips demand, by the rules of traditional economics the price of the commodity should fall to zero (this, for instance, can be demonstrated from the zero-profit theorem, Hirschleifer, Glazer and Hirschleifer, 2005, Chapter 10). In multinational schemes, where there is no analog of the medieval potentate seeking profits, the existence of political pressure towards awarding permits in excess of actual emissions is a generic feature of regulatory commodities.

3. Behavior of carbon permits prices in 2005-2007

I analyze two first contiguous years of the price history (2005-2007) for the three instruments: spot prices for emission credits, emission credit future for December 2007 and emission credit future for December 2008. Investigation of futures with tenor beyond 2008 did not produce significant new information with respect to the 2008 contract. Our sample ends on November 30, so as to coincide with the end of trading history for the first futures contract.

Price behavior of the three instruments is plotted in Figure 1 (see Appendix). We observe a sharp decline in prices of both the spot and the December 2007 contract in mid-2006. This decline was related to the fact that the market recognized that the EU had set a ceiling for emissions so high that there was little possibility that EU members were likely to never break through their country limits before the new reset of ceilings in 2008 (EU, 2007; Ellerman and Joshkow, 2008). Hence, polluters and their agents exited the market. The only remaining players in spot and 2007 futures contracts were the speculators gambling on the unlikely possibility that the limits would be breached. In a more rigorous language, information asymmetry between the regulator and the potential polluter always result in a number of outsiders willing to bet on an unlikely outcome (e.g., Hirschleifer, Glazer and Hirschleifer, 2005). This price collapse is viewed by some researchers as a signature of the intrinsic inefficiency of the emissions trading market (Victor and Cullenward, 2007).

The above reasoning already points at an option-like pricing mechanism of the commodity. The same pattern can be visualized in the terms of convenience yields (Hull, 1997) for the December 2007 and 2008 futures in Figure 2 (see Appendix). We define convenience yield by the standard equation:

$$q = -\frac{1}{T} \log \left(\frac{F e^{-rT}}{S} \right), \quad (1)$$

where F is a future price, S is a spot price, r is a risk-free rate and T is the time to maturity. For the brief discussion of convenience yields, see Appendix.

Superficial volatility for the 2007 future following the decline of the price is undoubtedly caused by rounding off the market prices to the nearest cent. Consequently, we plot 16-day moving averages for the convenience yield to remove the influence of the round off. Extensive statistical properties of my sample are listed in Table 1 (see Appendix). We used 1-year and 2-year swap rates as proxies for the risk-free rates for the December 2007 and December 2008 futures, respectively.

4. Model of the stochastic process for the spot price

I will start from a microeconomic description of the model. The above presentation is heuristic and is not based on any quantitative reasoning. A representative customer (a polluting firm) is able to pay a maximum fixed price for a unit of pollutant emission. This price depends on the available technology and profitability of the market and is generally unobserved by the regulator. The firm's business requires a certain quantity of permits, which is limited from above by Q^* . The regulator observes the demand and issues a certain number of permits. The regulating agency is a profit maximizing entity. If it supplies too few permits, they will deliver a maximum price but the quantity sold will be suboptimal. On the other hand, if it supplies too many of them, it can sell Q^* permits but the price per permit will fall sharply. A supply-demand curve reflecting the above reasoning is plotted in Figure 3 (see Appendix).

As one can observe from Figure 3, the model contains a single optimum slope of the supply curve, i.e. elasticity of response to the imbalance between supply and demand on the market of pollution credits. The regulator is supposed to maximize her implied utility, i.e. the implied profit from selling permits. In practice, the regulator (EU) distributes most of the permits for free, leaving the bulk of the implied profits on the table for political reasons (Axelrod, Vandever, Downie, 2010). However, one can hypothesize that even in the absence of the real cost of permit, the maximum potential profit indicates maximum efficiency of the regulation.

The stylized microeconomic model, that I propose above, has only zero or infinite elasticity of demand. This feature is obviously unrealistic but I consider it a good approximation in the short term. Indeed, in the short term, emissions reduction strategies requiring firms to undertake costly capital projects are out of question.

The supply-demand equilibrium formulated above is obviously devoid of dynamics. To develop a dynamic model allowing for demand shocks, we need a certain process that describes the approach to equilibrium. I took the imbalance and price processes as both being AR (1). The equation for the imbalance reads as:

$$u_t = \gamma U_0 + (1-\gamma)u_{t-1} + \sigma_1 e_t, \quad (2)$$

where u_t is imbalance at a time t , U_0 is an equilibrium demand for permits projected by the regulator, σ_1 is the volatility of the business model and $e_t \sim N(0, 1)$ ¹.

¹ Note that, because the price reaction to imbalances is non-linear (see next equation), the constant U_0 cannot be eliminated through the change of variables.

Parameter γ regulates the speed of the approach of equilibrium.

Price evolution is described by the system of equations:

$$p_t = (1 + \beta)p_{t-1} + f(u_t) + \sigma_2 \varepsilon_t, \tag{3}$$

$$f(u_t) = \begin{cases} 0, & u_t < 0, \\ \alpha u_t, & 0 \leq u_t < p_0 / \alpha \\ 0, & u_t \geq p_0 / \alpha \end{cases}$$

In equations (3), p_0 is the equilibrium price, α is the regulator's elasticity of supply, σ_2 is the volatility of external shocks to price, $\varepsilon_t \sim N(0, 1)$ and β is a parameter.

The intuitive meaning of equation (3) is that, if imbalance is negative – the firm does not need permits – the price of a permit in the next auction *drops* by $100 \times \beta$ percent in the absence of external shock. The same is true if the regulator issues permits in excess of production demands. Otherwise, the price reaction to the changes in the supply of permits is proportional to the imbalance. For my preliminary estimation (see next Section), I exclude the exogenous shock to the prices by setting $\sigma_2 = 0$.

$$\rho(\theta_0 | X(\theta)) = \frac{\text{Cov}[X(\theta_0), X]}{\sqrt{\text{Var}[X(\theta_0)]} \sqrt{\text{Var}[X]}} = \frac{\hat{E}[X(\theta_0)] - \hat{E}[X(\theta_0)], X - E[X]}{\sqrt{\hat{E}[(X(\theta_0) - \hat{E}[X(\theta_0)])^2]} \sqrt{E[(X - E[X])^2]}} \tag{4}$$

where the expectation sign with a hat means an average over computer-generated paths, while the expectation sign without hat means an empirical sample average.

Our estimation method is based on the fact that maximizing an absolute value of a correlation coefficient

$$I_{01} = E \left[\log \left(\frac{f_1(x, y)}{f_0(x, y)} \right) \right] = \int \log \left(\frac{f_1(x, y)}{f_0(x, y)} \right) f_1(x, y) dx dy. \tag{5}$$

The intuitive meaning of the KLD is the information we obtain if, instead of expected $f_0(x, y)$, the observed distribution is $f_1(x, y)$. Granger and Lin (1994) proposed a normalized entropy measure:

$$e_{01}^2 = 1 - \exp(-2I_{01}). \tag{6}$$

Normalized entropy (6) has obvious properties resulting from the properties of the KLD (for the proof, see Y. Hong, 2006):

1. $e_{01} = 0$ only if $f_0(x, y) = f_1(x, y)$;
2. $e_{01} = 1$ only if y is functionally dependent on x ;
3. e_{01} is invariant under transformation $x' = h_1(x)$,

exogenous shock not related to supply-demand fluctuations is sometimes called “microstructure noise” in the parlance of the microstructure theory. The likelihood of simulated paths is rather good even without adding it to the model.

The system of equations (2) and (3) describes the evolution of prices and volumes of capacity trading. The price is a non-linear function of permit imbalance.

5. The estimation of parameters for the spot prices from the 2005-2007 time series

I do not know the exact solution of the model formulated in Section 3. Thus, any guess of the maximum likelihood estimator is impossible. To estimate the parameters, I use the method of Bayesian inference (Tsay, 2002). Imagine that we have a model that generates an empirical sample $X(\theta)$ for a set of true parameters θ . For a given value of parameter vector θ_0 , one can generate a set of the model paths. If one postulates the criteria for the likelihood of the model path and the empirical trajectory of the asset, in this case, a correlation coefficient, one may form an expression:

cient maximizes the entropy. Standard logic runs as follows. Application of entropy in financial time series is a description of the affinity of distributions (Y. Hong, 2006). This application is based on the notion of Kullback-Leibler Distance (KLD, e.g., Lawler, 2006) between two alternative distributions: $f_0(x, y)$, which we consider baseline and $f_1(x, y)$:

- $y' = h_2(y)$, where h_1 and h_2 are smooth monotonic functions;
4. If $f_1(x, y, \rho)$ is Gaussian and $f_0 = f_1(x, y, \rho = 0)$, then $e_{01} = |\rho|$.

My optimization procedure in this Section is based on property (4). Namely, if one has a random sample X_i with parameters θ_i , then one can use normalized entropy as a sample distribution function of parameter values. Optimizing this distribution, according to Granger and Lin (1994), will be equivalent to maximizing entropy. My chosen method has a computational advantage because rather than extracting distributions from the time series, then calculating the integral (5) and, finally, maximizing I_{01} ,

one can compute the correlation between empirical sample and the simulated sample and then weigh observations according to $|\rho|$.

As usual with Bayesian methods, we must select a

$$\hat{\theta} = \underset{\theta_0}{\operatorname{argmax}} \left[\frac{\int \rho(\theta_0 | \theta) P(\theta) d\theta}{\iint \rho(\theta_0 | \theta) P(\theta) d\theta_0 d\theta} \right] \approx \frac{\int \theta \rho(\theta_0 | \theta) P(\theta) d\theta}{\iint \rho(\theta_0 | \theta) P(\theta) d\theta_0 d\theta} \cdot (7)$$

The logic behind equation (7) is that we estimate a true vector of parameters as maximizing the posterior distribution of the correlation coefficient between the empirical sample and the model-generated paths. The first expression is “exact” to the degree the simulated Bayes formula is exact; the second expression replaces the mode of the distribution with the median and is much more computationally parsimonious.

To compute the integrals (7) for several adjustable parameters (5=6, in our case), we use a suitable version of the Monte Carlo Markov Chain/Bayes procedure (Tsay, 2002). First, I generate a random number of paths sampled from a relatively arbitrary prior. Then, I compute the correlation of these paths with a true path. The values of parameters with their own correlation measure are weighted according to the Bayes law to obtain the posterior. In principle, this procedure can be iterated an arbitrary number of times, using the posterior of the previous stage as the prior in the new iteration in imitation of the Metropolis algorithm (Tsay, 2002). Yet, the first sampling from the gamma distribution (which involves 0.5 million simulated day-prices) already provides statistically significant answers.

Of course, this estimation method can work only if the model, being tested, produces paths that sometimes resemble actual empirical trajectories. Otherwise, the expressions of equation (7) will be random and numerically small. Note that once the parameter vector is estimated, Monte Carlo methods can be used to price any derivative on the spot or future contracts. We shall discuss the derivatives pricing in Section 7.

The results of my Bayesian estimation are given in Table 2. The comparison of the acceptable – highly correlated – simulated path with the empirical spot price is provided in Figure 4 (see Appendix). I must emphasize that the particular values of the parameters do not matter much - what is important at this stage is the principal possibility of building a model that describes collapses and revivals of the permit prices.

$$c(K, t, T) = \sum_{X \in \Omega} M[t, T] \cdot \operatorname{Max}[X - K, 0] \cdot \rho(X, \hat{\theta}) \approx \sum_t M[t, T] \cdot \operatorname{Max}[X - K, 0] \cdot \delta(X_t - X), \tag{8}$$

$$p(K, t, T) = - \sum_{X \in \Omega} M[t, T] \cdot \operatorname{Min}[K - X, 0] \cdot \rho(X, \hat{\theta}) \approx \sum_t M[t, T] \cdot \operatorname{Max}[K - X, 0] \cdot \delta(X_t - X),$$

prior $P(\theta)$, which can be any reasonable distribution. In this case, I choose gamma distribution. Now we can estimate a parameter vector from the Bayes formula:

6. Off-sample testing of the model

As the next test for the model, I generated another Monte Carlo panel with the parameters taken from Table 2 (see Appendix), and compared the first four moments of the empirical distribution of the prices. My sample this time used permits from the 2005-2008 sample, i.e. the sample in which I mix my 2005-2007 results from the start-up trading with about a year-and-a-half of data from more mature markets. The results of my comparison are provided in Table 3 (see Appendix). One sees that, despite an artificially agreeable form of the best-fitted entropic distribution (Figure 4, correlation coefficient ≥ 0.8), the third and fourth moments are reproduced rather poorly. In fact, the theoretical distribution plotted on the same scale is much more localized than the empirical one (Figure 5). The first idea that comes to mind is that shocks to the $AR(1)$ sequence, which models supply and demand, are not i.i.d. normally distributed. Our example shows that the entropic estimates, while providing time series qualitatively similar in shape, can err very significantly when the higher moments of the distribution are concerned. I must caution, though, that empirically we have two- or three-humped distributions. With conventional bell-shaped curves the situation can be different.

7. Additional insight from the options pricing

As we already have mentioned in the previous Section, the spot price model allows us to price derivatives. This exercise is far from academic because this model is non-linear and, hence, scales of the parameters’ change from dimensionless units into euros and numbers of permits (i.e. tons of carbon) matter. When the spot price process is established, pricing of calls and puts by Markov Chain Monte Carlo (MCMC) is straightforward. Call (put) valuation is based on the direct computation of the area above (below) a representative random path (Figure 6, see Appendix). We simply use a posterior distribution from equation (4) to calculate the expected payoffs:

where the second sum is taken over a substantial number of simulated trajectories (day-prices). In equations (8), $c(K,t,T)$ and $p(K,t,T)$ are the prices of call and put, respectively T is the expiration date (December 2008 in all our simulations); K is the option exercise price and $\hat{\theta}$ is the estimated parameters' vector. The factor $M[t,T]$ is a numeraire indicating conventional transformation into risk-free gauge. The risk-free rate for the past was taken from the Euro swap rates but for proprietary option pricing, it must be modelled, of course, through standard procedures (e.g., James, Webber, 2000 and Xie, Liu and Wu, 2005). The second equality is based on the assumption of ergodicity on the sample definition of a distribution function (Tsay, 2002).

$$p(X) = \sum_t \delta(X_t - X). \quad (9)$$

The option pricing method through MCMC is schematically shown in Figure 6 (see Appendix). Once the prices are converted into a riskless frame, a call (put) price in the same frame is simply the mean area above (below) a simulated spot trajectory. Furthermore, if one believes in ergodicity, one can use a sufficiently long path simulated with estimated parameters (Table 2) instead of a large number of paths in the first set of equalities in equation (7).

Estimated parameters for put and call options are provided in Table 4 (see Appendix). We notice that the pricing of calls from our postulated spot price process of equations (2) and (3) is quite consistent in the sense that all calls require the same scaling parameter independent of the strike price. However, in the pricing of puts not only scaling parameter does differ between the strikes, it is also an order-of-magnitude larger than the scaling parameter for the calls. This is highly suspicious given that upper-end strike prices for the puts overlap with the lower-end strike prices for the calls. Theoretically, this could happen if there were large violations of the put-call parity but the futures price, implied by the put-call parity for the options, gives quite reasonable prices (Table 5, see Appendix). These considerations seem to invalidate the model until we attempt to use a historical price distribution to value the same options. I use a special bootstrapping procedure to create a sample of exactly the same size as the simulated path ($T = 2^{13} = 8192$).

The procedure, which is technically closer to "jack-knife" than to "bootstrap" according to the original terminology but now is frequently united under the same umbrella term (Efron and Tibshirani, 1994; Shao and Tu, 1996). Namely, I randomly shuffle the daily price data for two years in my sample ($T = 519$), remove 7 prices from the reshuffled empirical series

and continue this procedure 16 times. After that, I append all 16 series in a single "long" trajectory, which has an equal length with simulated paths ($T_s = 8192$). As the case was with many numerical estimators, they have a bell-shaped maximum but fat tails. While we can use a longer trajectory, the probability of visitation of the tails becomes higher and the precision of the estimation by the t-criterion does not grow with the number of simulated observations beyond $T \approx 10^4$. Hence, I limited myself to sixteen replications.

Not only the call pricing based on historical prices is accurate, but also the scaling parameter q is of order unity, as it should be because historical and future prices have the same unit of measurement.

Puts priced by the bootstrapped historical price simulations demonstrate the same anomaly: the scaling factor is large and it experiences large differences between strikes. Hence, our invented price process shares both realistic and unrealistic features with the pricing based on bootstrapping the historical prices. We shall discuss possible reasons for that in the conclusion.

Conclusion

The proposed model of emission permit issuance demonstrates that price collapses and revivals do not necessarily indicate the failure of the regulatory market. In fact, they are an integral part of this market. A technical reason for this is that the price for a commodity, which can be created and withdrawn by a regulator's fiat, is a highly non-linear function of demand. The fundamental reason, though, is contained in the fact that the customer of a regulatory commodity has an incentive to hide from a regulator the quantity she needs and the maximum price she is willing to pay. From her point of view the payments for permits, similar to taxes, constitute a deadweight loss to the government, which provides no benefit to the producing firm. Price collapses, from the point of view of a producer, constitute an optimum state of affairs when the net cost of fulfilling the regulation is minimal.

Because of the sharply non-linear response of prices to the fluctuations in demand, the overall scaling of the simulated prices is important for practical valuations. I include a discourse on the subject, using additional information from the available options prices. There is no overall scaling factor that allows simultaneous correct pricing of calls and puts. Our model correctly prices only the calls. In the absence of observed violations of the put-call parity this could be viewed as a factor invalidating the model. Yet, the use of bootstrapped historical prices reveals

the same anomaly. The price of puts implied by the bootstrapped empirical distribution is several times higher than the market price. Independent of a particular pricing algorithm, we can expect that if the price of the commodity has a reasonable probability to fall to near-zero levels, the market price of the American put must hover near the strike price (Hull, 1997). Yet, the put has 30-40% of the price implied by historical spot prices, as well as our simulation.

I must note that the limited liquidity of the permits cannot explain this underpricing because one can create a synthetic put shorting a covered call. Thus, if the put is severely underpriced because of liquidity

concerns, so must be the call. However, this possibility is refuted by quite sensible December 2008 futures price implied by the empirical put-call parity (Table 5, see Appendix).

An amazing hypothetical possibility is that the depressed put prices are a direct manifestation of the free distribution of permits by the ETS. This distribution is equivalent in principle to the fiat money distributed by the EU governments (type I bubble in terminology of Jarrow, Protter and Shimbo, 2007; though the authors considered only a positive price bubble). Empirical methods to prove or disprove this hypothesis (Lerner, 2007) are still in their nascent stages.

References

1. Axelrod, R., S. Vandever and D.L. Downie (2010). The global environment: institutions, law and policy, CQ Press: Thousand Oaks, CA.
2. Basawa, I.V. and B.L.S. Prakasa Rao (1980). Statistical inference for stochastic processes, academic press, London, UK.
3. Cambridge (1987). Miller E., C. Postan and M. Postan (eds.) Cambridge economic history of Europe, vols. 1-3, Cambridge University Press: Cambridge, UK.
4. Çetin, U. and M. Verschuere (2008). Pricing and hedging in carbon emissions markets, *International Journal of Theoretical and Applied Finance*, 12, 7, pp. 949-967.
5. Daskalakis, G., D. Psychoyios and R.N. Markellos (2009). Modeling CO₂ emission allowance prices and derivatives: evidence from the European trading scheme, *Journal of Banking and Finance*, 33, 7, pp. 1230-1241.
6. Doyle, W. (1984). The price of offices in pre-revolutionary France, *The Historical Journal*, 27, 4, pp. 831-860.
7. Duffie, D. and K.J. Singleton (2003). Credit risk: pricing, measurement and management, Princeton University Press: Princeton, NJ.
8. Duby, G. (1974). The early growth of European economy, Cornell University Press, Ithaca, NY.
9. Efron, B. and R.J. Tibshirani (1994). An introduction to bootstrap, Chapman & Hall/CRC: Boca Raton, FL.
10. Ellerman, A. and P. Joshkov (2008). The European Union emissions trading system in perspective, Pew Center in Global Climate Research, May Report, Arlington, VA.
11. EU (2007). <http://ec.europa.eu/environment/climat/emission.htm>.
12. Granger, C. and J.L. Lin (1994). Using the mutual information coefficient to identify lags in non-linear models, *Journal of Time Series Analysis*, 15, pp. 371-384.
13. Hirschleifer, J., A. Glazer and D. Hirschleifer (2005). Price theory and applications: decisions, markets and information, Cambridge University Press: New York, NY.
14. Hong, Y. (2006). Nonlinear time series analysis: econometric theory and applications, Cornell University Lecture Notes: Ithaca, NY.
15. Hull, J. (1997). *Options, Futures and other derivatives*, Prentice Hall, Saddle River, NJ.
16. James, J., N. Webber (2000). Interest rate modeling: financial engineering, Wiley: New York, NY.
17. Jarrow, R.A., P. Protter and A. Shimbo (2007). Asset price bubbles in complete markets, pp. 97-122, in M.C. Fu (ed.), *Advances in Mathematical Finance*, Springer: Berlin, FRG.
18. Lawler, G.F. (2006). Introduction to stochastic processes, Chapman & Hall /CRC: Boca Raton, FL.
19. Lerner, P. (2007). Was there a bubble in NASDAQ?, Information-based re-investigation, working paper, www.ssrn.com/abstract=1150510.
20. Lerner, P. (2009). Microstructure and noise in financial markets, VDM Verlag: Saarbrücken, Germany.
21. Liptser, R.S. and A. N. Shyraev (1977). Statistics of random processes, Springer-Verlag: New York, NY.
22. Maland, D. (1966). Europe in the seventeenth century, A&C Black Publishers: London, UK.
23. Marshall, B.V. (1975). Comprehensive economics: institutional, analytical and applied, Longman: London, UK.
24. Mas-Colell A., M. Whinston and J. Green (1995). Microeconomic theory, Oxford University Press: Oxford, UK.
25. Mayers, C. (1957). The sale of peerages in early Stuart England, *Journ. Mod. Hist.*, 29 (1), pp. 21-37.
26. Merton, R.C. (1992). Continuous-time finance, Wiley-Blackwell: New York, NY.
27. Paolletta, M. S. and L. Taschini (2008). An econometric analysis of emission allowance prices, *Journal of Banking and Finance*, 32, 10, pp. 2022-2032.
28. PWM (2002). Appeals court upholds FERC decision in New England installed capacity case, *Power Markets Week*, Oct. 14 (14).
29. PWM (2005). FERC rejects complaints against ISO-NE's capacity, will address issue broadly, *Power Markets Week*, Apr. 11 (14).

30. PWM (2007). New York capacity market participants battle over legitimacy of bids and mitigation, Nov. 26 (1), *ibid.* New York ISO defends reform proposal for hotly disputed NYC capacity market, *Power Markets Week*, Dec. 17 (14).
31. Shao, J. and D. Tu (1996). *The jackknife and the bootstrap*, Springer: New York, NY.
32. Stavins, R. (1998). What can we learn from the grand policy experiment? Lessons from SO₂ trading, *Journ. of Econ. Perspectives*, 12 (3), pp. 69-88.
33. Tsay, R. (2002). *Analysis of financial time series*, Wiley, New York, NY.
34. Uhrig-Homburg, M. and M. Wagner (2009). Futures price dynamics of CO₂ emission allowance – an empirical analysis of the trial period, *The Journal of Derivatives*, 17 (2), pp. 73-88.
35. Victor, D. and D. Cullenward (2007). Making carbon markets work, *Scientific American*, 297 (6), pp. 70-79.
36. Xie, A., S. Liu and C. Wu (2005). Duration, default and the term structure of interest rates, *Journal of Financial Research*, 28, pp. 539-554.

Appendix. Approaches to the pricing of the pollution allowances

While most of the extant literature is dedicated to policy regulations and the economics of environmental mitigation (e.g., Axelrod, Vandever and Downie, 2010), the attempts to produce pricing method for pollution allowances are much fewer in number. In this Section, I review four such approaches restating their advantages and shortcomings. First is the paper by Paoletta and Taschini (2008) from the Swiss ETH. They analyzed time series for CO₂ and SO₂ emission credits using modified forms of GARCH in a purely econometric, reduced-form approach (for GARCH and its variants, e.g., Tsay, 2002). Their work produced a number of sharp observations, such as “spot-forward parity is inadequate due to the inconsistent behavior of the CO₂ emission allowance convenience yield, which depends on the political uncertainty that largely affect long term maturities.” The approach of Paoletta and Taschini can also be used for pricing derivatives. Yet, it does not include any substantive, no matter how stylized, economic model of the market for pollution credits.

Another approach, also based on econometrics of the time series, was undertaken by Uhrig-Homburg and Wagner in the series papers (e.g., Uhrig-Homburg and Wagner, 2009). They extensively studied the relationship between spot and futures markets. They suggest that after initial inefficiency, EU emission futures were rationally priced. In their observation, the futures markets led price discovery, which is quite obvious viewing the dramatic divergence between 2007 and 2008 futures prices approximately a year before reallocation of quotas (Figure 1) but requires quite tedious econometric research to resolutely prove it. As applied to the pricing of derivatives, they argued that EU allowances are traded and storable commodities. For the first, trial period of 2005-2007, they observed that the above statement is empirically correct and the standard risk-neutral valuation works. However, after the reallocation (2008), “the valuation of such derivatives should not be based on the current spot price because it does not reflect all the information necessary for building an expectation about future spot prices in the years 2008 and beyond. However, the future does reflect this information”. Indeed, politics of allocation of permits already settled down when the trading started in 2005 and could not affect spot prices and short-term futures and forwards. However, all the political considerations involved in initial allocation of the permits (national lobbying power, softer limits for the new members of the EU) must be replayed anew in the 2008 stage of allocation. Uhrig-Homburg suggests “pricing respective derivatives relative to the future 2008”, with relationship to which risk-neutral valuation is again applicable.

One complete pricing model is present in the paper by Daskalakis, Psychoyios and Markellos (2009). They estimate an affine jump-diffusion process and then price derivatives by the modified Black-Scholes-Merton approach (Merton, 1992). They describe the difference in 2007 and 2008 pricing by convenience yields. Yet, the information provided by the convenience yields, as we have seen, can be unstable (Figure 2) and deceptive (see above). The accuracy of their derivative pricing is $\pm 20\%$ in absolute magnitude and $\pm 10\%$ in percentages. The advantage of the jump-diffusion estimation method is that the practitioners used it for a long time and many proprietary-quality computer algorithms are available. Hence, it is likely to be the most useful for practitioners. However, this method is not specific with respect to the environmental commodities and does not imply anything in particular about the economics of permit trading.

The model, which presumes the existence of two intrinsic states (“high demand”-“low demand”) for a system, governed by an unspecified statistical process θ was developed by Çetin and Verschuere (2009). I also use the two-state paradigm, only in my case, the switching between states is performed endogenously through the explicit definition of market imbalance rather than exogenously, by a separate (unobservable) stochastic process. Their model belongs to a wider class of Markov-switching models (Tsay, 2002). Çetin and Verschuere develop equations, which govern the evolution of θ in a vector Markov model that can be empirically calibrated. The equations (4.10) and (4.11) of their work are similar to the two-state filtering problem described by Liptser and Shyraev (1977, Chapter 15.4). The approach of Çetin and Verschuere is so far the most sophisticated mathematically but is also the most complicated for the purposes of empirical estimation. Not for nothing, unlike the three above-cited articles and the author’s paper, which calibrate the market data, Çetin and Verschuere provide examples with stylized, but reasonable values of parameters. It will be interesting, for instance, to calibrate a pricing model in

their approach and then compare the dynamics of their unobservable “anxiety” parameter θ and the real market demand.

Table 1. Descriptive statistics of the prices in EU emissions of two-year time series

Parameters	Spot	Future '07	Future '08
Mean	9.955	10.460	19.894
Standard error	0.427	0.451	0.157
Median	8.600	8.800	19.500
Mode	0.070	0.070	18.250
Standard deviation	9.738	10.285	3.566
Kurtosis	-1.189	-1.169	0.597
Skewness	0.470	0.485	0.680
Range	29.700	31.440	20.000
Minimum	0.050	0.060	12.250
Maximum	29.750	31.500	32.250
Confidence level(95%)	0.840	0.887	0.308

Note: Descriptive statistics of the prices in EU emissions permits from 11/28/2005 to 11/27/2007. Distributions of spot prices and the December 2007 future can be made nearly identical by insignificant scaling. The distribution of December 2008 future prices is very different and corresponds to the expectation of permit re-allocation some time in 2008.

Table 2. Estimates of the model parameters

Model parameter	Expected value	Standard deviation
α	0.00306	0.0017*
β	-0.0559	0.0290
γ	0.0303	0.0198
ρ_0	0.317	0.169*
U_0	0.568	0.310*
σ	0.0627	0.0309

Note: Estimates of the model parameters in equations (3) and (4) of the main text obtained by the method of Bayesian inference (Basawa and Prakasa Rao, 1980; Tsay, 2002). The standard deviation of parameters valid at 5% is listed in boldface; the parameters with an asterisk are valid at 10%.

Table 3. Comparison of the moments of the price distribution

Moments	Empirical distribution	Theoretical distribution
Mean/median	1.460	1.676
Mean/std	0.939	0.866
Skewness	0.856	2.232
Kurtosis	-0.318	7.358

Note: Comparison of the moments of the price distribution obtained from the theory of Sections 2-5 and the distribution of the complete sample of 2005-2008 prices. The model with AR(1) demand provides reasonable relationships of the mean to the median and to the standard deviation (first two columns) but the third and fourth momentum diverge very sharply.

Table 4. Bootstrapped historical and MCMC simulations of the option prices expiring in December 2008

Strike	MCMC simulation			Bootstrapped historical prices		Option
	q_{sc}	Avg. (2007) price, €	Mispricing, €	q_{sc} /implied strike	Mispricing,€	
€10	1088	0.76	0.25	75.7	0.18	put
€12.50*	854	1.26	0.17	79.0	0.25	
€14.50*	512	2.15	-0.33	58.6	-0.17	
€15	613	2.04	-0.13	69.0	0.05	
€20	426	3.95	-0.93	47.8	-0.62	
€21.50*	607	3.62	-0.24	70.4	0.12	
Average 683.4				66.7		
€20	82.8	4.64	0.10	1.54/€13.0	1.23	call
€22.50*	88.0	4.60	-0.53	1.66/€13.6	0.24	
€25	82.0	3.31	0.16	1.61/€15.5	0.57	
€30	82.2	2.41	0.12	1.67/€18.0	0.10	

Table 4 (cont.). Bootstrapped historical and MCMC simulations of the option prices expiring in December 2008

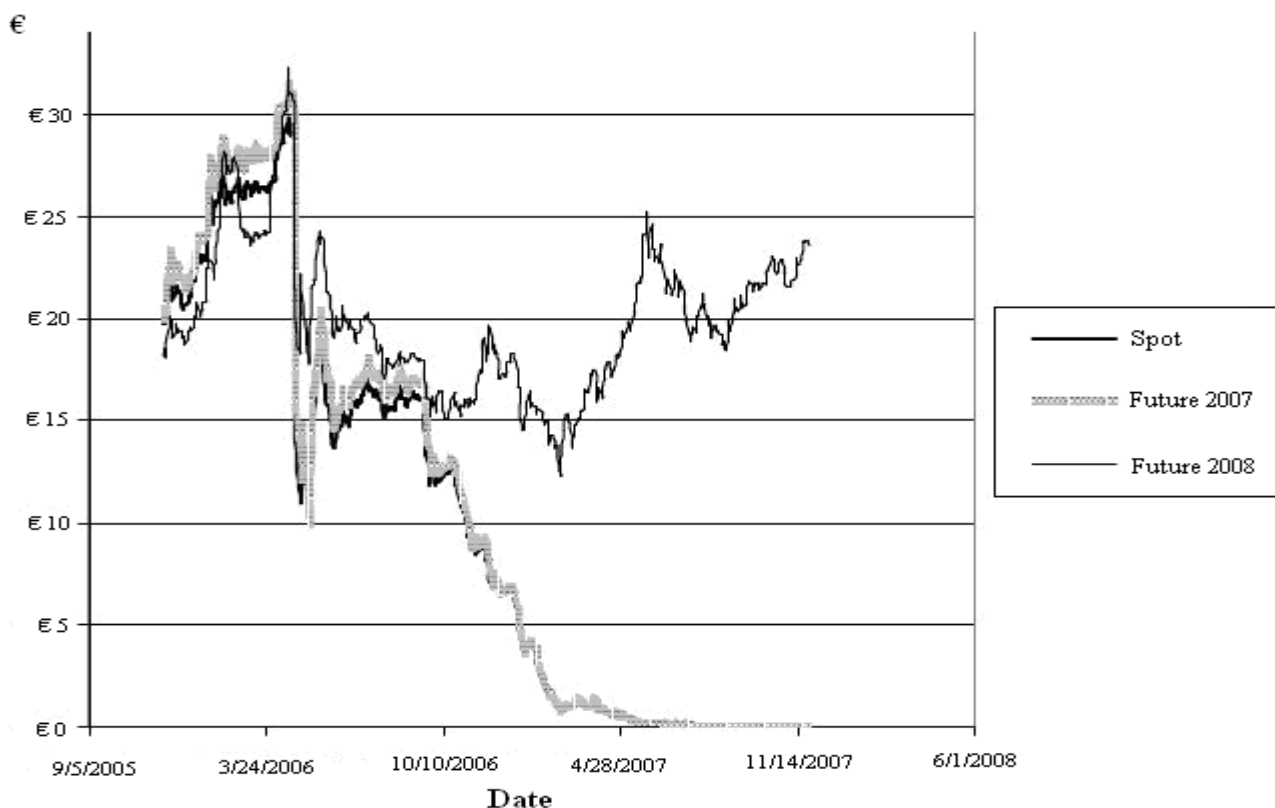
Strike	MCMC simulation			Bootstrapped historical prices		
	q_{sc}	Avg. (2007) price, €	Mispricing, €	$q_{sc}/\text{implied strike}$	Mispricing, €	Option
€35	91.1	1.45	-0.56	1.87/€18.7	-0.93	
€40	75.1	1.02	0.40	1.78/€22.5	-0.33	
Average 83.5				1.69		

Note: Exact prices can be matched with a scaling factor q_{sc} for arbitrary continuous distribution – a consequence of a theorem from elementary calculus. To calibrate the model, we use mispricing of our valuation procedure with respect to the average price of Dec-2008 options when using an average scaling factor for puts or calls only. For some strike prices (indicated by asterisk) synthetic options were created by averaging options with close strikes. The a priori scaling factor for the historical simulations must be 1. Calls approximately correspond to this rule. Excess of the actual strike price with respect to the strike price implied by simulation is approximately equal to the ratio of mean and median spot price. In fact, spot prices above €30 were not observed for the entire trading sample (2005-2007) but the calls with strikes up to €50 were written during the period.

Table 5. Prices of futures in 2007

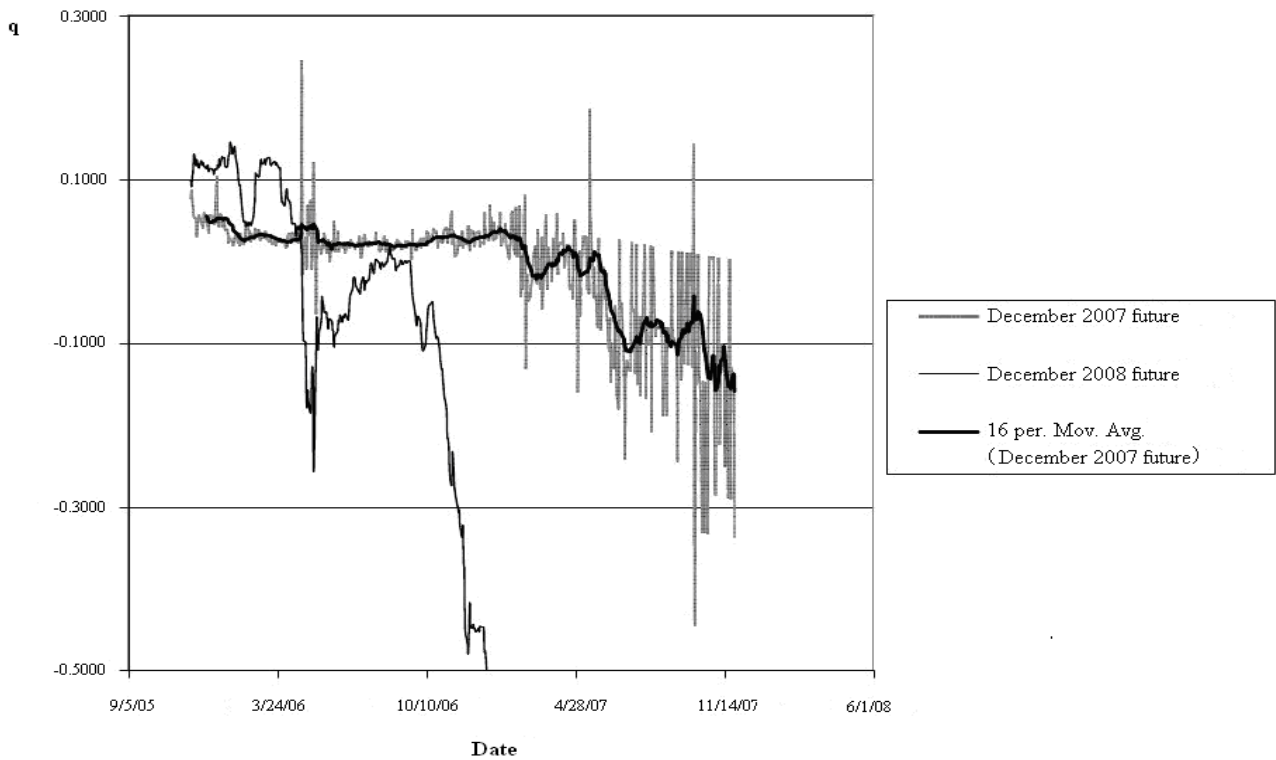
Month	Price, €
June-07	21.59
July-07	20.42
Aug-07	19.42
Nov-07	22.49

Note: Price of futures in December 2008 implied by a put-call parity. All futures prices are given on the end-of-month basis. For other months of 2007 there is an insufficient number of puts and calls with the same strike prices.



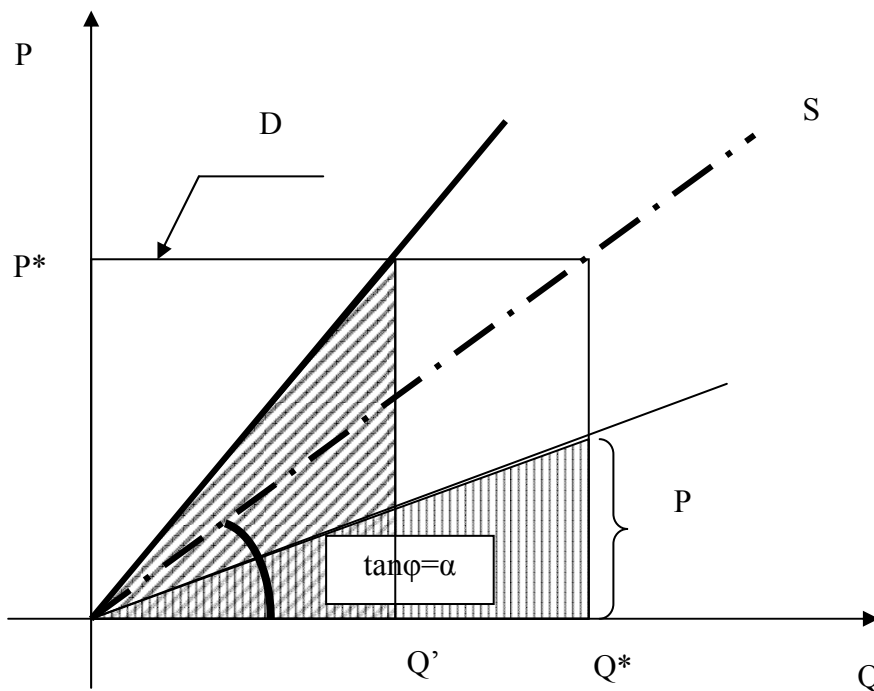
Note: Spot prices are plotted in boldfaced black, the future expiring in December 2007 in smudged grey and the future with expiration in December 2008 in solid black. The December 2007 future follows the spot price, while the December 2008 future sharply deviates from their dynamics two years before expected re-pricing

Fig. 1. Prices of EU emission permits in euros for the two year period 11/28/2005-11/27/2007



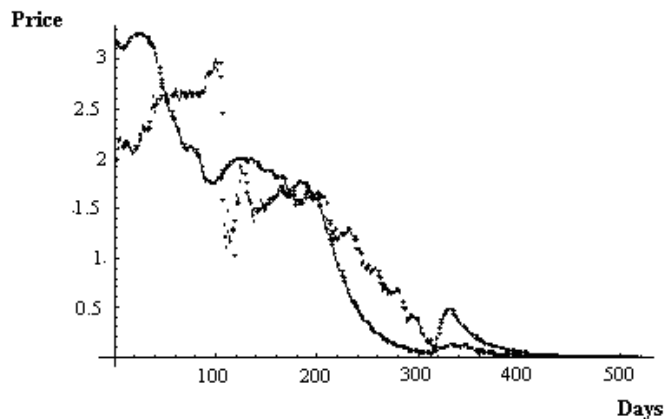
Note: Convenience yield becomes sharply negative two years before the expected re-pricing, which indicates the expectation that emission quotas will be reduced.

Fig. 2. Convenience yield of December 2007 future (smudged grey) and December 2008 future (solid line) and 16-day moving average for December 2007 future (boldfaced solid black)



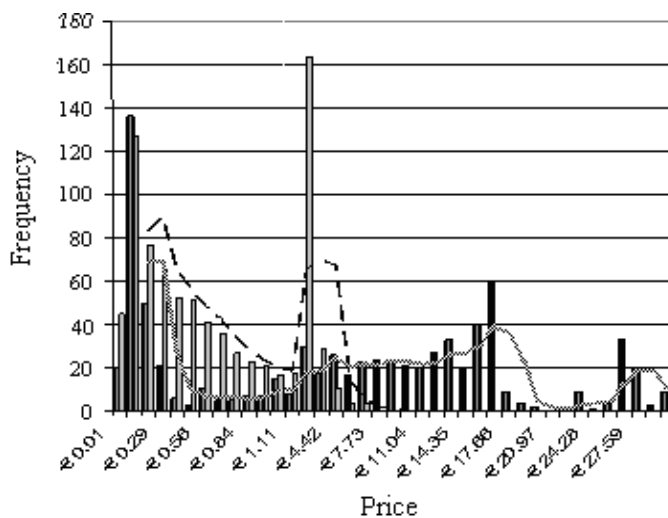
Note: Demand is constant at P^* for $0 < Q < Q^*$. Regulator selects α , the elasticity of the supply of permits. Oblique shading denotes the area that equals half of the profit for the regulator in the case of suboptimal quantity Q' , and the vertical shading – half-profit for the suboptimal price P . Maximum profit is achieved when the elasticity of supply corresponds to optimal quantity and optimal price.

Fig. 3. Idealized supply and demand curves for the regulatory commodity



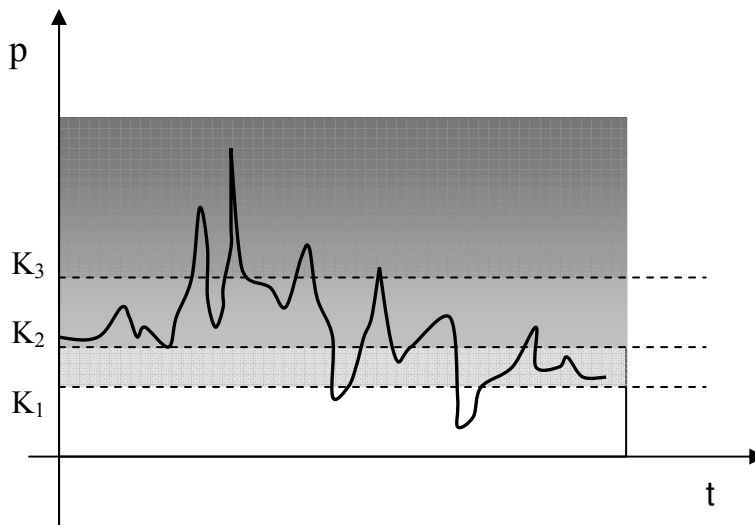
Note: The horizontal axis is plotted in units of trading days from the beginning of the sample (11/28/2005) two years into the sampling history. The vertical axis is plotted in units of the ratio between daily and median price.

Fig. 4. Comparison of one of the simulated price trajectories (dotted line) with the empirical price (dots)



Note: In Figure 5, a simulated curve has been scaled to conform to the maximum observed market price. The solid grey line and the dotted black line are the moving averages of 4 columns of each distribution provided as a guide for the eye. We see that while the distributions have qualitatively the same shape, under this choice of scaling factor, the theoretical distribution is skewed to the lower prices.

Fig. 5. Theoretical (black) and empirical (grey) prices for the model of Sections 2-5 and the 2005-2007 sample



Note: The price of call with the strike K_1 is the area under the curve and above K_1 (indicated by grey). The price of put is the area under the curve below K_1 (indicated by white). Pricing algorithm works for an arbitrary strike not necessarily traded on the market, even if the price path is empirical.

Fig. 6. Schematic demonstration of the derivatives pricing using simulated trajectory in a risk free frame