STRATHCLYDE

DISCUSSION PAPERS IN ECONOMICS

FORECASTING THE EUROPEAN CARBON MARKET

BY

GARY KOOP AND LISE TOLE

NO. 11-10

DEPARTMENT OF ECONOMICS
UNIVERSITY OF STRATHCLYDE
GLASGOW

Forecasting the European Carbon Market[1]

Gary Koop, Department of Economics
University of Strathclyde
&
Lise Tole[2]
University of Edinburgh Business School

February 2011

ABSTRACT: In an effort to meet its obligations under the Kyoto Protocol, in 2005 the European Union introduced a cap-and-trade scheme where mandated installations are allocated permits to emit CO_2. Financial markets have developed that allow companies to trade these carbon permits. For the EU to achieve reductions in CO_2 emissions at a minimum cost, it is necessary that companies make appropriate investments and policymakers design optimal policies. In an effort to clarify the workings of the carbon market, several recent papers have attempted to statistically model it. However, the European carbon market (EU ETS) has many institutional features that potentially impact on daily carbon prices (and associated financial futures). As a consequence, the carbon market has properties that are quite different from conventional financial assets traded in mature markets. In this paper, we use dynamic model averaging (DMA) in order to forecast in this newly-developing market. DMA is a recently-developed statistical method which has three advantages over conventional approaches. First, it allows the coefficients on the predictors in a forecasting model to change over time. Second, it allows for the entire forecasting model to change over time. Third, it surmounts statistical problems which arise from the large number of potential predictors that can explain carbon prices. Our empirical results indicate that there are both important policy and statistical benefits with our approach. Statistically, we present strong evidence that there is substantial turbulence and change in the EU ETS market, and that DMA can model these features and forecast accurately compared to conventional approaches. From a policy perspective, we discuss the relative and changing role of different price drivers in the EU ETS. Finally, we document the forecast performance of DMA and discuss how this relates to the efficiency and maturity of this market.

Keywords: Bayesian, carbon permit trading, financial markets, state space model, model averaging

[1] Financial support from the ESRC under grant RES-062-23-2646 is gratefully acknowledged. Both authors are Fellows at the Rimini Centre for Economic Analysis.

[2] Address for correspondence: Lise Tole, Edinburgh University Business School, 29 Buccleuch Place, Edinburgh, EH8 9JS, U.K. Email: lise.tole@ed.ac.uk.

1 Introduction

The largest carbon market in the world, the European Union Trading Scheme (EU ETS) commenced operation on January 1st, 2005. It is the key institution whereby EU countries aim to meet their obligations under the Kyoto Protocol. Currently, over 11,000 energy-intensive installations in the power industry and five major industrial sectors (including oil, iron and steel, cement, glass, and pulp and paper) are included in the scheme. Together these account for nearly half of Europe's total CO_2 emissions. The scheme caps CO_2 emissions, and then distributes allowances to emit to large energy users. One EU allowance (EUA) is equal to one metric tonne of CO_2. This "right to pollute" permit is a tradeable commodity. Thus, the supply of permits is set by the cap, and the demand for permits depends on the level of CO_2 emissions in a given year. In 2009, the total value of the market had grown to €103 billion by the end of the year, with 8.7 billion tonnes of $C0_2$ traded, accounting for over 95% of the total value of carbon traded worldwide (Kossoy and Ambrosi 2010).

The ultimate goal of the EU ETS is to finance investment in a low carbon economy. According to analysts, if this is to be met, the price of carbon must rise substantially to at least €40 per tonne. However, the all-time high was achieved in 2006, when prices reached just over €30 a tonne, but usually it has been much lower than this. Compounding the problem of low prices is their volatility. When trading began in 2005, the price of a permit on the spot market was around €5. It rose quickly to between €20 and €30 and then crashed on the announcement that 2005 allocations had been overestimated. Prices rebounded again to €15 before falling to under €1 in February 2007 and to less than a cent in August 2007. Prices for 2008-12 allowances began trading at over €20, before falling throughout the latter half of 2008. The market regained some ground in 2009 and, since the middle of 2009 has been fluctuating between €12 and €16.

Like any commodity market, certain fundamentals (often referred to as "price drivers") will be expected to play a role in explaining changes in the carbon price. But the exact list of price drivers and the magnitude of their affects on price is unclear. Moreover, unlike other commodity markets, the European carbon market is fundamentally determined by compliance. The supply of permits is set by the European Commission (EC), which also controls the allocation, bankability and size of the program through the number of participants and permits. These elements of the trading scheme have differed between the first (2005-2007) and the second trading phase (2008-2012). They will also change again, in the third phase (2013-2020), as the cap tightens and the size of the market increases with the extension to more polluters. Uncertainty over a global climate agreement, the size of the future cap and market, the eligibility and number of some types of offsets allowable in the scheme are key factors impacting on the price of carbon now and in the future. While all commodities markets face uncertainties, the changing conditions of imposed by compliance are unique to the EU ETS.

Because of the changing complexities of this commodity market, measuring price drivers and forecasting prices raises particular challenges for the statistician seeking to understand the dynamics of this new and evolving market. As with many financial studies (e.g. Avramov 2002), the number of potential variables that may affect prices is large. This can lead to over-parameterization problems (i.e. regression techniques may apparently fit well in-sample, but this may be due to over-fitting, leading to poor forecast performance). In response to this problem, it is increasingly common to use Bayesian model averaging in such cases (see e.g. Avramov, 2002, Ley and Steel, 2009 and 2010 and the references cited therein). Moreover, the marginal effects of predictors can change over time, i.e. parameters can change or structural breaks can occur. Also, the relevant forecasting model may change. These types of changes are especially likely to occur in a market such as the EU ETS, where there are many potential macroeconomic, financial and institutional variables, the impacts of which will vary both within and between phases, with implications for the price of carbon. To address these issues, we use an approach called Dynamic Model Averaging (DMA). DMA can handle the over-fitting problems caused by the presence of many predictors and allows for both parameter and model change and, thus, seems ideal for forecasting the carbon markets.

We forecast spot and future prices of carbon permits traded between April 2005 and August 2010 using a set of large set of potential predictors which combines those used in a variety of other studies. We find that DMA forecasts of carbon prices in the EU ETS are much better than conventional forecasts. Furthermore, DMA gives insight into the range and changing role of various factors driving carbon prices in the different phases of the EU ETS.

The remainder of this paper is organized as follows: Section 2 provides an overview of the EU ETS market since 2005, its main institutional aspects and price drivers. It includes an overview of the empirical literature on carbon trading as it relates to the EU ETS. This section also motivates the choice of variables used in the analysis. Section 3 discusses definitional and other issues relating to the data set. Section 4 presents the DMA methodology. Section 5 presents the paper's results. Section 6 concludes.

2 Carbon Prices in the EU ETS

2.1 Institutional and Policy Aspects

In contrast to the command-and-control approach to pollution control, which sets and then regulates limits on emissions, the EU ETS uses a market-based approach.[3] Under the Kyoto Protocol, the EU has agreed to cut emissions by 8% relative to 1990 levels by the years 2008-12. This may rise to as much as 30% cuts relative to 1990 levels by the end of 2020. Under the EU ETS scheme, large emitters subject to compliance are allocated an annual number of permits. Each permit grants the right to emit one tonne of CO_2. Permits can be bought and sold in financial markets. An installation that exceeds its cap can purchase

[3] More detailed information on the EU ETS can be found at http://ec.europa.eu/clima/policies/ets/index_en.htm

additional permits on the market. An installation that has more permits than emissions can sell its excess permits or bank them for a limited period. At the end of each year, companies must surrender a number of permits equal to the amount of their total emissions for the year.[4] Each year the EC, the body responsible for regulating the scheme, discloses information on compliance and verified emissions for each installation and at the country level.[5] Emissions and compliance information on all member states' installations under the scheme is contained in the Community Independent Transaction Log (CITL), which is publicly available and updated each year.[6]

Emission allowances are allocated to installations for a number of phases, known as trading periods. The scheme has three trading periods, defined by their different compliance requirements. The first trading period (Phase 1) ran from 2005-2007. During this phase, nearly 2.2 billion permits of 1 tonne each were allocated each year. Approximately 60% of these were allocated to the power sector. Allocation is based on each member state's National Allocation Plan (NAP), which in turn is based on its Kyoto commitment. The NAPs were used to decide how many allowances are allocated and to which installations. NAPs have been criticized for leading members to misrepresent their emissions, leading to an overallocation of permits. As a consequence, caps were tightened for Phase 2 (2008-2012), which are 6.5% lower than in Phase 1. The vast majority of permits have been allocated freely to installations, based on their historical emissions.[7] Member states were permitted to auction up to 5% of allowances in Phase 1 and 10% in Phase 2. Limitations on the auctioning of permits have been criticized for suppressing carbon prices and creating windfall profits for installations.[8]

In order to provide flexibility for installations, the scheme allows for intertemporal borrowing and banking of permits within each phase. That is, firms can borrow on next year's permits to fill a shortage in the current year or bank permits for upcoming years. Inter-phase borrowing and banking was not

[4] However, they are not required to hold the number of permits during the year; a set-up that raises the possibility of liquidity problems if installations find that they are suddenly short of permits when it it is time to surrender them.

[5] Specifically, the calendar of compliance is as follows: On February 28th of each year installations receive their new allocation for the upcoming commitment year. March 31 is the deadline for the submission of a verified emissions report. April 30th of each year is the deadline for surrender of permits used in the previous year. May 15 is the deadline for the report by the European Commission of verified emissions for all installations covered by the EU-ETS.

[6] The CITL can be viewed at http://ec.europa.eu/environment/ets/.

[7] There is also a requirement that new entrants receive a reserve of free allocations while allocations from entities that have shut-down cannot be re-sold. Ellerman and Buchner (2008) discusses the implications of these institutional features and finds their impacts on prices to be ambiguous.

[8] As a consequence, critics argue for the auctioning of permits as a way of efficiently allocating permits and ensuring a higher price. Auctioning also provides other benefits, including fewer tax distortions, greater incentives for businesses to make investments in low-carbon technologies, and the reduction of political manipulation in the allocation of permits (Cramton and Kerr 2002; Bohringer and Lange 2005).

allowable in Phase 1 or Phase 2. However, allowances distributed during Phase 2 can be banked for future use in Phase 3. Banking/borrowing provisions may have an impact on the carbon price. Banking may help installations hoping to hedge risks against seasonal and cyclical price swings. However, on the other hand, interperiod restrictions may severely affect the price of carbon. Alberola and Chevallier (2009b), for example, demonstrate how the restriction in banking and borrowing between Phase 1 and Phase 2 was responsible for dramatic price falls observed in 2007.

One notable feature of Phase 2 has been the introduction of carbon offsets. As an alternative to obtaining EU ETS carbon permits, a firm may offset some of its carbon emissions by investing in emission reductions elsewhere in the world. The mechanism for doing so was set up through EU's Directive 2004/101/EC (The Linking Directive), which links the EU ETS with project mechanisms under the Kyoto Protocol.[9] Under this directive, firms can purchase a limited number of carbon offsets called CERs (certified emission reductions) and ERUs (emission reduction units) in order to meet their obligations.

The fundamental economic rationale of this scheme is that countries differ in their marginal costs of abatement. Poorer countries, in particular, will tend to have lower marginal abatement costs than industrialized countries of the EU. Hence these mechanisms help to achieve a given emission reduction target at a lower cost than possible at home. In 2008, CERs accounted for 3.9% of all surrendered permits, the overwhelming number of these originating from projects in China and India. ERUs accounted for only 0.002% of surrendered permits. Thus, in Phase 2, carbon offsets accounted for a small, but important part of the EU carbon market. Trading in CERs began in earnest in 2008 while ERUs only recently started trading in the latter half of 2010. In Phase 2, installations have been able to buy international credits equivalent to 1.4 billon tonnes CO_2 or a yearly average of 280 million tonnes per year. In Phase 3, rules on the use of credits is still undecided but installations will be able to carry over unused credits from Phase 2 and an undetermined number of additional offsets. In any event, the use of offsets will be limited to no more than 50% of the EU emissions reductions to be made between 2008-2020 and after 2013 only offsets from countries that have approved an international climate agreement or new types of EC approved projects will be permissible.

The CER market itself has many interesting institutional features. CER prices opened in the first months of January 2008 at over €20, with prices falling to below €10 on both the spot and futures markets in February 2009. Overall, prices have averaged around €13 since trading began. Despite this relative stability, project financing for CERs has diminished, falling over 50% between 2009 and 2010. A significant event affecting the price of CERs was the announcement by the Hungarian government on March 11, 2010 that it

[9] These mechanisms are called the Joint Implementation (JI) and the Clean Development Mechanism (CDM). The CDM allows for investment in projects that reduce or offset emissions in developing countries; while the JI is limited to projects in countries that have agreed to cuts under Kyoto, i.e. industrialized and transitional market economies. CERs are generated by the CDM and ERUs by the JI.

had sold nearly 2 million CERs which had been surrendered for compliance to its national register. Surrendered permits are automatically retired in the EU ETS, but similar rules were not in place for CERs. Hence there was nothing stopping the Hungarian government from selling the CER permits it was holding back on the market. The news that recycled permits were circulating on the international market led to a fall in the price of CERs.

Another problem to hit the CER market centred on the questionability of its projects. The vast majority of CDM investments have been in China and involve the destruction of hydroflourocarbons (HFCs) and nitrous oxides, which are potent global warming gases. They accounted for about a half of the 430 million CERs issued up to 2010 (Point Carbon 2010), helping to create a certain price stability for CERs. When news emerged that many projects had involved expanding production of HFCs simply to gain CERs by shutting them down again, prices fell by nearly 2%. The CDM executive responded by delaying the issuance of CERs from such projects, a move that increased CER prices. As a consequence of these problems, from 2013, offset credits from certain HFCs and nitrous oxides will likely be banned.[10]

Fraudulent activity has also emerged in the EU ETS itself. It came to light in the beginning of 2009, when high volumes of trade were noted on France's Bluenext exchange. Trading reached a record 19.8 million permits on 2 June, 2009, having exceeded the previous record of 15.1 million tons on 29 May, 2009 (Frunza et al. 2010). This heightened activity was due to so-called "carousel fraud". The fraud involves criminals buying permits free from VAT in one EU country, and then selling them in another country with VAT added but without paying the VAT to the relevant country tax authority. Europol estimated the total cost of the VAT fraud to be around €50 billion (Frunza et al. 2010).

"Phishing attacks" constitute another form of fraud affecting the market, the first occurring in early February 2009. The first attack saw an estimated 250,000 permits with a value of over €3 million stolen by criminals setting up fake registries. The attack involved six German companies giving their registration details away. It led to the temporary shut down of the trading system on February 2nd (BBC 2010). A second attack in early January 2011 led to the theft of €28 million from national registries in Austria, Greece, the Czech Republic, Poland and Estonia. As a consequence, spot trading was shut down in all national registries for several days (New York Times 2011).

Due to the need to ensure that increasingly stringent emissions targets and ensure a high and stable price for carbon, the EC has implemented a number of changes for Phase 3 (2013-2018). These measures include: a lengthening of the compliance phase to 8 years (2013-2020) and an end to national allocation plans for individual member states. These plans will be replaced by a centralized EU-wide cap on emissions, which will include all greenhouse gases and all sectors. This cap is designed to eliminate the biases of the previous system, wherein

[10] For the world as whole, issuances of CERs fell to 132 million tons in 2009, 10% less than in 2008. In the first quarter of 2010, only 10 million CERs had been issued. The EU ETS will no doubt be affected by this shortage. Forecasts suggest that demand for the certificates is expected to fall by 1 billion tonnes by 2012 (Kossoy and Ambrosi 2010).

member states favoured their own industries. The cap for 2013 will be based on the average total quantity of allowances in Phase 2 and then a reduction factor of 1.74% will be applied linearly each subsequent year. In Phase 3 there will also be an increase in the number of permits auctioned to a minimum of 50% (with 100% auctioning in the power sector in most of the EU). As mentioned, the use of offsets is likely to be severely restricted. These and other measures are expected to lead to a 21% reduction below 2005 verified emissions by 2020.[11] If an international agreement on climate change is reached post-Kyoto, the cap will be tightened at 30% below 1990 levels.

At the beginning of 2011, uncertainty casts a shadow over the EU ETS market. Economic recovery remains fragile in Europe and the euro is under threat. Several carbon schemes in other parts of the world have floundered and the Copenhagen and Cancun climate conferences have failed to reach a global agreement on reducing greenhouse gas emissions. Regulatory uncertainty surrounds how CERs and ERUs will be credited in the next phase and what quantitative limits will be placed on their use. It is also uncertain which CERs will be tradable (i.e. whether HFC projects are banned). Uncertainty also surrounds the impact on the carbon market of the EU's renewable energy targets, which have been set at 20% of total energy use by 2020.[12]

2.2 Empirical Studies of Price Drivers in the EU ETS Market

The previous sub-section discussed political and institutional issues that have influenced the carbon market. The general picture is one of change, as policymakers fine-tune the rules of the game so as to improve the system, and as various outside pressures (e.g. fraud) influence the markets. We now turn to a discussion of how the existing literature has attempted to statistically model this changing and turbulent market.

An important focus of existing studies is the role of energy markets and associated events (e.g. weather) in influencing carbon prices. Since industrial and power sectors have received more than 50% of EUAs, this focus is understandable, as their behaviour largely determines demand (and thus prices) on

[11] The allocation of allowances to individual sectors will depend on benchmarks based on the average of the top 10% most CO_2 efficient installations in the EU. The problem of carbon leakage (i.e. industries relocating to countries outside the EU so as to avoid reducing carbon emissions) will be addressed by giving sectors with a significant risk of locating overseas 100% of the benchmarked allocation for free. Those less likely to move will receive 80% of their benchmarked allocation for free in 2013, which will fall to 30% in 2020 and 0% in 2027. The 300 million allowances from the new entrants reserve of the EU ETS will support the demonstration of carbon capture and storage (CCS) and innovative renewable technologies.

[12] In their 5th survey of participant sentiment in the GHG market, the International Emissions Trading Association (IETA) found that 2/3rds of respondents believed that current uncertainties will have a significant negative impact on long-term low-carbon investment. Confidence in the EU-ETS and the CDM was also down on the previous year's survey. One-half of respondents predicted that CERs will no longer be the dominant international offset by 2015.

the market.

The energy switching behaviour of power operators is of particular interest, which depends on the options available to them. Specifically, power installations in the EU generate electric power largely using natural gas and coal. To meet their emission limits operators have a number of options: They can reduce the amount of electricity they produce, but this rarely happens. They can move toward low-carbon installations, which is the explicit aim of the EU ETS, but hardly a short-term option. Only when the price of carbon rises high enough to replace coal-fired generation with low carbon alternatives will this be an option. Consequently, in the short-term, the usual option open to generators is to take advantage of the abatement opportunities provided by switching from coal to gas-fired generation. Coal tends to be cheaper than gas but it also emits more than twice the CO_2 emissions per mwh of electricity produced. Operators will tend to consider the difference in relative profits from electricity generation using coal and gas. Simply by switching to gas, they can achieve some abatement of their CO_2 emissions, and as a consequence, will require fewer permits. Economic theory states that in equilibrium, the price of a permit should equal the marginal cost of abatement. In the short-term, abatement will be based on switching from coal to gas. Thus the marginal abatement cost should depend on the difference between the price of gas and the price of coal. If gas prices rise, the marginal abatement cost will also rise and thus the price of a permit should as well. However, if coal prices rise such that coal generation is only a bit cheaper than gas generation, marginal abatement costs will fall, as will the price of a permit. In short, the price of coal and gas are important factors in driving carbon prices, through the ability of producers to engage in fuel switching behaviour (Christiansen et al. 2005; Chevallier 2009).

Empirical studies have indeed found associations between energy prices and the price of carbon. Keppler and Mansanent-Batallet (2009) demonstrate that during Phase 1, coal and gas prices, through clean dark and clean spark spreads[13], impacted on CO_2 future prices. Similarly, Alberola et al. (2008b), Hintermann (2010) find that high gas prices are associated with higher carbon permit prices. They also find that, in a regression seeking to explain the carbon price, coal prices and the clean dark spread are significant and of the expected negative sign.

Bunn and Fezzi (2007) find the carbon price reacts quickly to the gas price. Kanen (2006), Convery and Redmond (2007) and Mansanet-Batallet et al. (2010) also find that energy price changes are significantly related to carbon prices, with increases in the price of oil (which prices of natural gas are closely tied to) having the most substantial effect.

Energy demand is also influenced by weather conditions. For instance, extremes of temperature in either direction (i.e. cold winters or hot summers) could lead to unexpectedly high demand for energy. This would cause emissions

[13] These spreads are the margin a plant makes from selling a unit of electricity, having bought the required fuel and carbon permits. The dark spread refers to coal and the spark spread to gas. The key point to note is that they heavily reflect the price of coal or gas, respectively.

to rise and thus affect the price of a carbon permit. Mansanet-Batallet et al. (2010) use as explanatory variables EU wide fuel prices and a weather index of several cities in the first year of the carbon market and find that the temperature in Germany is the only significant driver of carbon permit prices. Hintermann (2010) finds that temperatures affect the price of carbon more significantly in the 2006 post-crash period in permit prices. Nordic reservoir levels and precipitation in Nordic countries (a proxy for reservoir levels in these areas) are also found to be negative and significantly related to carbon prices. However, their effect is diminishing, which reflects the fact that reservoir capacity is limited as is the amount of rainfall captured. Alberola et al. (2008b) focus their weather analysis on extreme weather events, i.e. hot summers and cold winters in 4 countries (Great Britain, France, Germany, Spain), weighted by the share of each country's NAP. The authors find that unanticipated extremely cold (but not hot) weather events are significantly and positively related to changes in carbon prices.

Other market fundamentals, such as macroeconomic and industrial measures, have been found to be empirically associated with carbon prices. Chevalier (2009) examines the role of commonly examined macroeconomic drivers in driving futures price volatility (i.e. equity dividend yields, the junk bond premium, the US Treasury bills, and the excess return on a global commodity index). The author finds that macroeconomic factors,with the exception of the default yield spread in the bond market and the dividend yield in the equity market, explain only some of the variance in carbon prices and only at certain points in time. Bredin and Muckley (2010) measure the impact on carbon futures of equity prices and an index of industrial production. Both of these variables are significant in the two phases, although industrial production was found to have a counter-intuitively negative sign.

Examining the role of sectoral output in the combustion, paper and iron sectors in large allowance-holding countries (France, Germany, Italy, Poland the UK), Alberola et al. (2009a) find that output in the combustion sector in countries (except in France and Italy) significantly impacted on EUA price changes in Phase 1. But, the coefficient on output in the iron sector had a negative sign, which could be explained by the net long position of this sector (i.e. too many permits were allocated to this sector). A variable controlling for a structural break from 25 April to 23 June 2006, when prices plunged on the announcement that many countries were over allocated permits, was significant for all countries with the exception of Germany. The authors attribute the latter finding to the net long position of this country's sector in this period. Alberola et al. (2008a) also find that only production in combustion and iron sectors in the EU 24 were significant explanatory variables in regressions involving EUA price changes in Phase 1.

In addition to the above market fundamentals, institutional features and events have been modelled for their potential role in affecting the price of carbon. Typically, these features are measured by the use of dummy variables. We will discuss such dummy variables in the next section.

3 Data

3.1 General Discussion of Properties of the Data

Our dependent variables are based on the spot and future prices of a carbon permit. The future has settlement date at the end of 2012 (i.e. this is the futures contract with settlement at the most distant period available in our data set). Our daily data set runs from 22 April, 2005 through 18 August, 2010.

Most empirical studies of the carbon market use either the spot price or the futures price. As we shall discuss below, both prices are of interest. They shed light on different features of the carbon market (e.g. futures will have a tendency to reflect long-run expectations about the state of the carbon market, whereas spot prices are potentially more sensitive to price drivers which have a short-run impact) and have different statistical properties. For these reasons, our study uses both.

Figure 1 plots the spot and future prices of a carbon permit. Two important aspects of the spot price series can immediately be seen. The first is evident in the middle of 2006, when the spot price dropped dramatically from roughly €30 to €10. This occurred in response to the first release of emissions accounting data for the EU ETS. This emissions verification showed that too many permits had been released in the sense that actual emissions in 2005 were well below the number of allocated permits. Subsequent to the price collapse in April 2006, a brief stabilization of the price took place until late 2006. But by early 2007 the spot price had fallen again to less than €1, with a carbon permit becoming virtually worthless (10 euro cents or less) by mid 2007.

The second feature is the sudden jump in the spot price, which occurs at the beginning of 2008, when Phase 1 of the EU ETS ended and Phase 2 began. This rapid jump in price occurred with the switch from worthless Phase 1 permits to Phase 2 permits. Recall that EU ETS carbon permits from Phase 1 could not be banked for use in Phase 2. Thus, formally speaking, in the case of the spot price the product being bought is a Phase 1 carbon permit for use in 2005-2007 and in Phase 2, a carbon permit for use in 2008-2012. For our futures price series the product being bought is always a Phase 2 carbon permit.

In order to address some of these data features, various statistical procedures are used in the literature. Some papers carry out statistical tests for structural breaks (e.g. Chevalier 2009), use estimation methods designed to be robust to structural breaks, work with sub-samples of the data, or use other methods. As described above, others use knowledge of institutional detail of the sort described in Section 2.1 to create dummy variables. For instance, Hintermann (2010) uses an "emissions verification dummy" which equals one on 25-28 April of each year (which is the time early information was released in 2006 about the previous year's emissions). Alberola et al. (2009b) and Chevalier (2009) use a dummy variable for the entire period, 25 April through 26 July, 2006. Alberola et al. (2009b) additionally use a dummy variable for 30 March through 30 April of each year (a period which includes the yearly verification announcements). Alberola et al. (2008b) include dummy variables for April 2006 and also for the period after October 2006 when a second slump in the carbon price occurs. Frunza et al. (2010) argues that the carousel fraud described in Section 2.1 had a substantial impact on carbon prices and specifies the period 1 November, 2008 through 31 August, 2009 as the time period the fraudsters were active, using a dummy variable for this period. Similarly, Mansanet-Bataller et al. (2010) use a dummy variable for the credit crisis (covering the period 17 August, 2007 through 31 March, 2009) and four dummy variables that relate to news that could impact on the carbon market concerning the release of information about National Allocation Plans in Phase 2, developments in the EU ETS plans relating to Phase 3, and the meetings of various executive boards and other groups for making plans relating to the CER market and the linking of emission trading schemes worldwide. In short, a myriad of different dummy variables have been used by different authors.

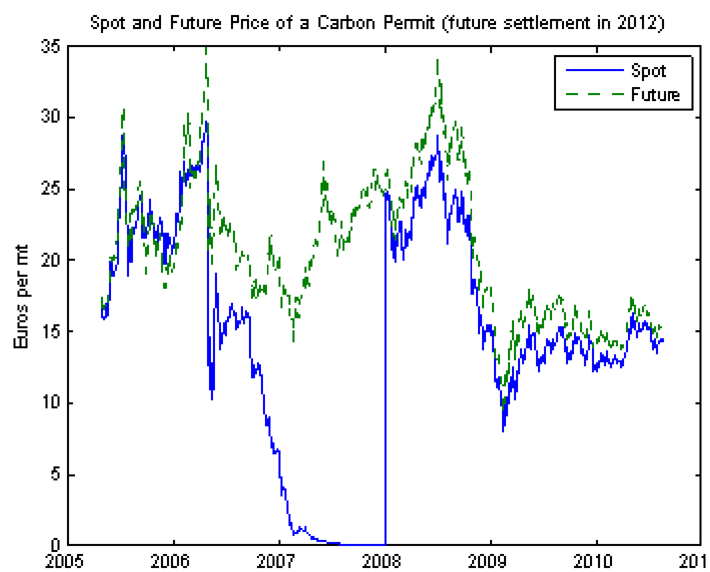

Figure 1

Does the EU ETS provide a suitable market for forecasting or addressing important questions of interest (e.g. what are the price drivers for carbon prices?)? As discussed in Section 2.1, strong evidence exists of repeated structural change in the EU ETS carbon market due to its various institutional features. For instance, it is possible that the carousel fraud problem in May/June 2009 or the price collapse in late 2006 that carried through to 2007 caused by the overallocation of permits make these periods so different from each other (and from other times) that it is meaningless to attempt a statistical analysis which includes them all. That is, it is possible that these factors completely explain the time series dynamics in these periods and that there is no role for the price drivers to play. It also may be the case that Phase 1 and Phase 2 are so different from one another that it is meaningless to pool them both together into a single statistical analysis. But we would expect this not to be the case. Even though institutional details and unique aspects of the carbon market undoubtedly play an important role in carbon pricing, it is likely that the price drivers discussed in Section 2.1 as related to persistent market fundamentals, also play a role, even though their effects may be weakened or masked by institutional features and events. The challenge, taken up in this paper, is to devise a statistical methodology that can deal with this and the other issues raised above.

The existing literature exhibits a belief that it is possible to answer important financial questions with carbon market data and attempts to address the issues noted above in various ways. One strategy in the literature is to work with futures instead of spot price data. From Figure 1, it can be seen that the futures price exhibits fewer of the pathologies that occur with the spot price. For instance, Bredin and Muckley (2010) use futures data and document the emergence of a stable market where prices are driven by fundamentals. However, the price of futures is not exactly the same as the spot price. Indeed several financial theories involve investigation of the relationship between spot and futures exchange rates. The price of a future reflects investors' beliefs about the price of a carbon permit at the settlement date. The further into the future the settlement date is, the more futures and spot price dynamics might differ. The future in Figure 1 has a settlement date of December 2012. The future price with yearly settlement dates (i.e. in each year the price is for the future with settlement in December of that year) yields a time series which is virtually the same as the spot price series. Thus, the future with yearly settlement exhibits all the pathologies of the spot price in Figure 1. It is only by considering settlement dates in the more distant future that we can remove these pathologies. But futures with such distant settlement dates can have a very different interpretation than the spot price.

As discussed, another strategy is to use dummy variables to control for institutional features. In the context of a forecasting exercise, a major problem with such an approach is that often these dummy variables are selected in retrospect. For instance, a forecaster in early April 2006 would probably not have known to include a dummy variable in late April. Even if she had, she would not have known its coefficient. Thus, forecasts cannot be based on dummy variables chosen in retrospect.

A second problem is that this strategy may not be sufficient. That is, the inclusion of dummy variables allows for the intercept of a regression to shift, but not the marginal effects of the price drivers.[14] In our case, it is plausible that the marginal effects will also change in response to the factors noted above. A third problem is that there are so many potential dummy variables, and the researcher has a degree of discretion in defining them. Consider for instance,

[14] Unless dummies are interacted with the price drivers. However, this strategy can lead to a proliferation of explanatory variables.

the emissions verification dummy commonly used in these studies to handle the institutional feature that, in late April of each year, an announcement is made about the actual emissions of the previous year. As noted above, some authors use a dummy variable for the period 25-28 April, while others use a dummy for the more extended time period from 30 March through 30 April. Some authors use such a dummy variable only for a period in 2006, but others include the emissions verification dummy variable each year (note that the latter may be inappropriate since the change in price on 25-28 April, 2006 is much larger than for other years). To handle the long price decline after the emissions verification of April 2006, some authors include a dummy variable for the entire period (25 April through 26 July, 2006), while others do not. Some authors include an additional dummy variable for the period around October 2006 when the second slump in the carbon price occurred. Similar uncertainties over the precise number and timing of dummy variables hold for other features of the data.

Thus, the existing literature does not have a systematic way of handling the changes that undoubtedly occur with this data set. It is common practice for the researcher to observe an odd pattern in the data and then to think of a story to justify that pattern. The trouble is that there are so many possible "stories". Consider for instance the story told about the decline in spot prices between April 2006 and the end of 2007. As we have seen, this decline is attributed to the fact that, in April 2006, the emissions verification process showed that the allocation of carbon permits was well above actual emissions in 2005. This undoubtedly does account for a large part of the patterns in the spot price. However, it does not explain other things (e.g. why the carbon permit price stabilized in the latter half of 2006).

Figure 2 plots allocated and actual verified emissions in each year. The over-allocation of permits in 2005 can clearly be seen. However, in other years, such as 2008 an under-allocation occurs (probably due to the increasing use of carbon offsets). But in 2009 there was an over-allocation of permits of the same order of magnitude as in 2005, but this did not cause a similar collapse in prices. All in all, great care must be taken in the story-telling used to create dummy variables in order to come up with sensible and consistent choices for such dummies.

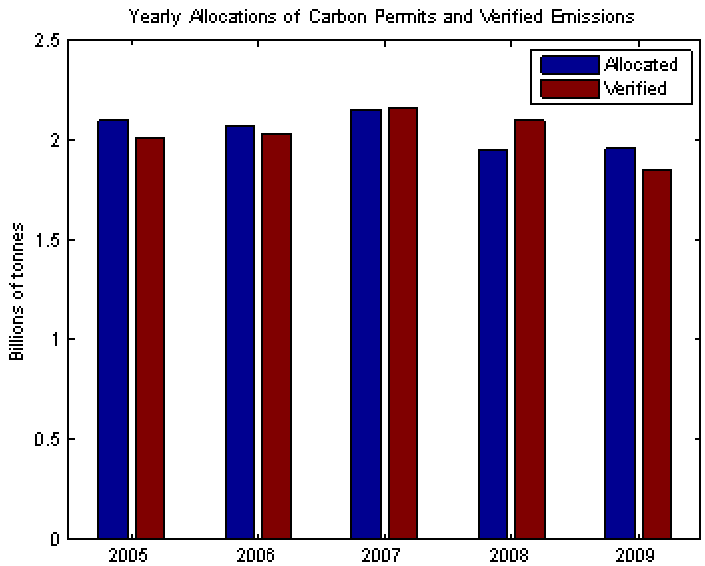

Figure 2

Recall that, in Phase 2, firms were able to use carbon offsets (instead of carbon permits) to partially cover their CO_2 emissions. For Phase 2, Figure 3 plots the CER futures price and the carbon permit futures price, both with a settlement data at the end of 2012. An examination of this graph suggests that, even though EU firms can only partially cover their CO_2 emissions by using carbon offsets, the two markets are clearly closely related to one another. This raises the question of whether the set of price drivers for an EU ETS carbon permit may be different in Phase 2 than Phase 1. That is, the price drivers in Phase 1 might reflect European considerations, whereas Phase 2 price drivers could be those affecting both the international CER market and the EU ETS market. For instance, unusually hot summers which increase electricity demand would increase the demand for carbon permits in Phase 1 and Phase 2. But in Phase 2 this driver might have less of an effect on the EU ETS price since firms could choose to buy carbon offsets on the CER market, which (since it is a world-wide market) would be less affected by an unusually hot European summer. This institutional change provides a further justification for a statistical methodology that allows for the marginal effects of price drivers on the carbon price to change over time.

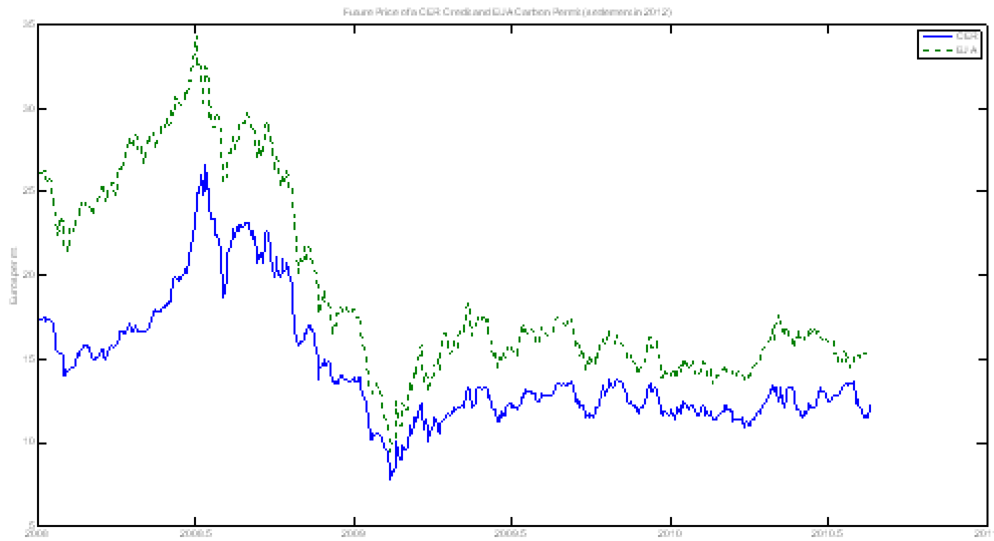

Figure 3

Another approach used by some papers for surmounting such data problems is to use only data for Phase 1 or Phase 2 (although papers such as Bredin and Muckley 2010, do use data that span both phases). However, the restriction of the analysis to one phase omits potentially useful information in both phases about financial and other features of interest, such as the role of the price drivers. Hence a more accurate estimation of such features of interest will be obtained from using all the data. However, it is important to use a statistical methodology that allows for parameters and models to change across phases.

3.2 Data Description

Before turning to our statistical methodology, we specify and justify our set of potential explanatory variables (or "price drivers"). These are chosen as covering a wide range of issues and have all been used before in the literature (see Section 2.2). For future reference, it is worth noting that this list of predictors is quite a long one, raising statistical concerns about over-fitting. A summary of data sources and acronyms is available in Appendix A.

a) **Dependent variables:**

The daily spot price of carbon permits (EUA) is taken from the Bluenext and Powernext exchanges. Data from the Bluenext Exchange, one of the most liquid spot market, was unavailable for the early months of 2005; hence we supplement the daily data for this period with data from the Powernext spot exchange. Measured in euros/tonne.

Daily futures price of carbon permit (EUA) for settlement on December 2012, the furthest date ahead for which we had complete data for both phases. As with the spot price, this variable is measured in euros/tonne of CO_2.

b) **Explanatory variables:**

We use the following independent variables reflecting a wide range of macroeconomic/financial and institutional variables (acronyms in brackets):

- Oil price (**poil**). Daily futures (month) ahead price of oil (brent in euros/barrel) from the Intercontinental Commodities Exchange (ICE). There is evidence to suggest that natural gas price movements in Europe are linked to those for crude oil (e.g. Asche et al. 2006 and Josse-Vasquez and Neumann 2006)

- Coal price (**pcoal**). McCloskey NW Europe Steam Coal marker.[15]

- Gas price (**pgas**). Measured as the Netherlands TTF day ahead continental gas futures negotiated in one of Europe's largest exchanges, APX-ENDEX. Measured in euros/mwh.[16]

- Electricity price (**pel**). Measured as the Phelix base load daily one month ahead price of electricity negotiated on the European Energy exchange in euros/mwh.

- Temperature (**tempdev**). The temperature of each EU country was calculated[17] and a weighted average across countries was taken (weights were proportional to population). The absolute value of the deviation from this mean temperature was calculated (see Bredin and Muckley 2010). Data were from the European Climate Assessment Dataset (Klein Tank et al. 2010). Note that carbon permit spot prices are expected to be most sensitive to unexpected increases in demand for electricity or oil. Such unexpected increases should be associated with unusually hot or cold days. Hence, the use of an absolute deviation from mean.

- Availability of hydropower energy variables (**resprop)** and (**precip**). The first of these variables was constructed from data from the Noordpool exchange which consists of the main hydropower producers in Europe (Norway, Sweden, Denmark and Finland). It was constructed by taking weekly reservoir capacity in the three countries and dividing by the total maximum possible. It was obtained from the Noordpool exchange. Reservoir capacity will affect the supply of hydroelectric power as will precipitation

[15] An alternative measure, which has been used in a number of studies (e.g. Alberola et al. 2009b), is the daily coal month ahead futures price. However, this variable is not as comprehensive in scope as the McCloskey marker as values are missing for one-third of the beginning of the time period under investigation

[16] This data set is more comprehensive than the commonly used alternative Zeebrugee day ahead measures for UK gas, which has missing data for 2005.

[17] A few small countries with patchy data were omitted.

(Hintermann 2010). Consequently, the second variable is the precipitation level in two Noordpool countries for which complete daily data were available, Sweden and Finland .

- Stock price (**pstock**). Daily index of the stock prices of the most highly capitalized, 100 blue-chip companies in Europe (the Euronext 100). This variable is a proxy for performance of the financial markets in the EU. It was derived from Thompson Datastream.

- Commodity price (**pcomm**). Daily index of world commodity prices (including energy). The Reuters commodity index measures the price of basic inputs into the European economy; these, in turn, will have an impact on the costs of production for installations and thus the demand for permits. This measure was derived from Thompson Datastream.

- Corporate risk premium variable (**junkprem**). The risk premium is an important component in the assessment of the value of corporate investments. A commonly used measure is the junk bond premium, calculated in this study as the difference between Moody's BAA and AAA-rated bonds. The higher the spread, the greater the credit risk. Other macroeconomic studies have included some interest rate measure as well, to indicate investment risk, information which would be incorporated in the above measure, which reflects this risk more directly.

- Carbon prices in the US (**pcarbus**). Daily carbon permit prices on the Chicago Climate Exchange (CCX). A voluntary carbon market in the US, the commodity traded is the CFI (carbon finance instrument). Participants voluntary agree to a legally binding commitment to cut emissions. The CCX is the largest carbon exchange in the US, and thus has the potential to affect the size of the market, if a comprehensive mandatory cap-and-trade legislation were ever to be passed. In the absence of mandatory legislation for now, a strong voluntary market in the US (strengthened by anticipation among traders of new mandatory legislation) may send a bullish signal to compliance markets elsewhere. The variable is measured in US dollars per metric tonne of carbon.

- Carbon offset price (**CERfuture**). We use the daily price of CER futures since 2008. We use the futures data rather than spot data on CERs since the latter has incomplete coverage. It is derived from the European Climate Exchange (ECX) and is measured as the daily CER futures price in euros with settlement in December 2012.

- Dummy controlling for the two phases **(dumphase1)**: Equals 1 during Phase 1 and 0 during Phase 2.

- Overallocation of permits measure (**overalloc**). This variable is constructed by subtracting allocated permits from actual emissions each year.

4 Statistical Methods

In this paper we are interested in forecasting the European carbon market. This interest derives partly because developing statistical methods for forecasting in such markets is of interest in and of itself (e.g. for investors in the market). But it is also important for the private sector to accurately assess future costs and plan necessary changes in production. Likewise, policymakers need such information to guide them in the setting of efficient emissions limits and for measuring the social costs of carbon-emission reduction. The pseudo real-time nature of a forecasting exercise allows us to see if there are times when more accurate forecasting is possible. In addition, a forecasting exercise is an excellent way of answering questions of the type: "What are the price drivers for the carbon market?" "Are these price drivers changing over time?". A forecasting exercise is less likely to suffer from over-fitting problems (i.e. a model with many potential price drivers could over-fit in-sample, with the over-fitting problem revealed in poor out-of-sample forecast performance).

The considerations of the previous sections mean that we want a statistical methodology that:

1. Is an extension of a regression model where a dependent variable (e.g. the price of a carbon permit) depends on a large set of explanatory variables.

2. Allows for the forecasting model to change over time.

3. Allows for the marginal effects of the regression model to change over time.

4. Avoids problems with over-fitting that can occur when the number of explanatory variables is large.

5. Allows for changes in volatility (which is an important consideration for any analysis involving daily financial data).

In this paper, we use a statistical methodology called DMA, developed in Raftery et al. (2010), which satisfies these characteristics. The reader is referred to Raftery et al. (2010) for complete details of DMA. Here we outline only the basic ideas, with some additional technical details provided in Appendix B.

To explain DMA, we begin with the time-varying parameter (TVP) regression model:

$$\begin{aligned} y_t &= Z_{t-1}\theta_t + \varepsilon_t \\ \theta_{t+1} &= \theta_t + \eta_t, \end{aligned} \tag{1}$$

where y_t is the dependent variable (e.g. the price of a carbon permit), Z_{t-1} is a $1 \times m$ vector of observations on explanatory variables that are used for forecasting y_t, θ_t is an $m \times 1$ vector of regression coefficients, ε_t is $N(0, H_t)$ and η_t is $N(0, Q_t)$. Note that we are forecasting one day ahead and, thus, Z_{t-1} is the information available for forecasting y_t (and in our empirical work, our price drivers will always be lagged one day). The TVP regression model is a state space model of the sort commonly used in empirical macroeconomics (see, e.g., among many others, Cogley and Sargent 2005; Cogley, Morozov and Sargent 2005; Primiceri 2005).[18] Standard methods (e.g. involving the Kalman filter) for estimation and prediction exist with such models. They allow for the marginal effects of the predictors to change over time. This is potentially very useful with the present data set. For instance, if the gradual decline of the carbon permit price in late 2006 was associated with a gradual change in the role of some of the price drivers, then a TVP regression model would be an ideal way of modelling this.

However, TVP regression models can deal poorly with the sort of abrupt changes in the time series dynamics that we may expect with this data set. Furthermore, they can be over-parameterized. That is, if m is moderately large (as in the present application) and the data are noisy, even a standard regression model can over-fit. In regression problems with many potential explanatory variables, Bayesian model averaging (BMA) has been a common response to such worries (see, among many others, Raftery, Madigan and Hoeting, 1997, Fernandez, Ley and Steel 2001a,b; Eicher, Papageorgiou and Raftery 2010 and Ley and Steel 2010). However, until very recently, it has been considered computationally infeasible to do BMA with TVP regressions such as (1). The contribution of Raftery et al. (2010) was to develop computationally feasible methods of doing model averaging with dynamic models such as (1). The resulting statistical methods are called DMA. The empirical work in Raftery et al. (2010) involves an engineering application, but Koop and Korobilis (2010) find DMA to work well in a macroeconomic application.

To define the set of models used in DMA, let $Z_t^{(k)}$ contain a subset of of the price drivers in Z_t for $k = 1, .., K$. Since Z_t contains m price drivers and there are 2^m subsets of these, we have $K = 2^m$ models. Each of the models in our model space is:

$$\begin{aligned} y_t &= Z_{t-1}^{(k)} \theta_t^{(k)} + \varepsilon_t^{(k)} \\ \theta_{t+1}^{(k)} &= \theta_t^{(k)} + \eta_t^{(k)}, \end{aligned} \tag{2}$$

$\varepsilon_t^{(k)}$ is $N\left(0, H_t^{(k)}\right)$ and $\eta_t^{(k)}$ is $N\left(0, Q_t^{(k)}\right)$. Let $L_t \in \{1, 2, .., K\}$ denote which model applies at each time period and $y^t = (y_1, .., y_t)'$.[19] The fact that we are

[18] Note that some of the carbon market literature uses cointegration methods. However, this is unnecessary here. When working with TVP models, researchers typically do not worry about unit root or cointegration issues. TVP models are very flexible and their focus is on fitting the time variation in coefficients. Even if variables have unit roots this is not a problem in the sense that all that matters is that ε_t and u_t do not have unit roots. The time-varying θ_t can effectively pick out cointegrating relationships (if they exist) leading to stationary residuals. Even if they fail to do so, note that the presence of a time-varying intercept which follows a random walk will pick up any remaining nonstationarity in the model.

[19] For notational simplicity, we do not explicitly include Z_{t-1} in the conditioning argument in the densities below (but all densities are conditional on the price drivers).

letting different models hold at each point in time and will carry out model averaging justifies the terminology "dynamic model averaging". To be precise, DMA involves estimating $\theta_t^{(k)}$ in each individual model and averaging these model-specific estimates across the model space using $\Pr\left(L_t = k|y^{t-1}\right)$ for $k = 1, .., K$ as weights. When forecasting, DMA involves averaging across predictive densities or point forecasts using $\Pr\left(L_t = k|y^{t-1}\right)$ for $k = 1, .., K$ as weights. Since standard Kalman filtering methods provide us with the predictive density in each model (see Appendix B), all that is required is a method for calculating $\Pr\left(L_t = k|y^{t-1}\right)$. Raftery et al. (2010) develops computationally efficient methods for this calculation.

The advantage of DMA is that it allows for switches between parsimonious models. That is, instead of suffering the over-fitting problems that often occur with TVP regression models that include all m price drivers at each point in time (or suffering from misspecification problems that arise if constant coefficients models or models with too few price drivers are used), DMA can switch from (say) a TVP model with three or four price drivers to a TVP model with three or four different price drivers. In practice (see, e.g., Koop and Korobilis 2010) DMA has been found to favour parsimonious models and can handle abrupt changes in the dynamic structure (such as might have happened when Phase 1 ended and Phase 2 begins) much better than conventional TVP models.

So far, we have said nothing about the error variances, $H_t^{(k)}$ and $Q_t^{(k)}$. Suffice it to note here that both of these will vary over time in our application. With financial data, it is especially important that $H_t^{(k)}$ be allowed to vary over time to capture the changes and clustering of volatility that characterize financial time series. We use an Exponentially Weighted Moving Average (EWMA) estimate of $H_t^{(k)}$ commonly used with financial time series (see, e.g., Riskmetrics 1996).

Further details on our implementation of DMA are given in Appendix B. Raftery et al. (2010) provides complete details and derivations. Note, in particular, that DMA involves the choice of two so-called forgetting factors, λ and α, which influence the rate of change in coefficients and models, respectively. These forgetting factors are set to $\lambda = \alpha = 0.99$ which are the choices recommended by Raftery et al. (2010).

5 Forecasting Carbon Prices in the EU ETS

5.1 Introduction

All the explanatory variables in our models are lagged by one day relative to the dependent variable. All of the models include an intercept, a lag of the dependent variable and a dummy variable which equals 1 for Phase 1. Our models differ in their treatment of the 13 price drivers given in Section 3.2. Thus, at each point in time, DMA is model averaging over $K = 2^{13} = 8192$ models (and the weights in the model averaging process are changing over time).

We divide our results in two parts: one applies to estimation using DMA; the other to forecasting. All variables are logged except for those which take on zero or negative values. We repeat all our empirical work twice: once using the log of the spot price as the dependent variable, the other using the log of the futures price as the dependent variable.

5.2 Price Drivers

Figures 4 through 9 summarize the information provided by our many models and parameters when the log of the carbon future is the dependent variable. Figure 4 sheds light on how parsimonious DMA is. Let $Size_{(k)}$ be the number of price drivers in model k then

$$E\left(Size_t\right) = \sum_{k=1}^{K} \Pr\left(L_t = k | y^{t-1}\right) Size_{(k)}$$

can be interpreted as the expected or average number of price drivers used in DMA at time t. Figures 5 through 8 present the inclusion probabilities for each variable (i.e. they are the weight used by DMA attached to models that include a particular price driver). Figure 9 plots the estimate of the error variance, H_t.

In general, Figures 4 through 8 show that, even with the carbon futures data (which does not have the same pathologies as the spot price data, see Figure 1), constant coefficient models that simply include all the price drivers are inappropriate. Figure 4 indicates that although 13 price drivers are available, DMA chooses to forecast with models with roughly half of them omitted. Figures 5 through 8 demonstrate that there is a great deal of variation over time in respect to which price drivers are included. And Figure 9 exhibits a great deal of time-variation in volatility.

With regard to Figures 5 through 8, we hesitate to tell strong stories about individual price drivers in the context of a reduced-form forecasting exercise such as the present one. Nevertheless, a few patterns are worth noting. It is rare for DMA to attach a probability close to one to any particular price driver at any point in time. This is to be expected with a financial forecasting exercise (where often the first lag of the dependent variable provides most of the predictive power) with correlated explanatory variables. As with previous studies, though, the prices of gas, oil, coal and electricity are often important price drivers. Note, in particular, the increasing role of natural gas prices in 2009.

In terms of financial variables, the junk bond premium (a measure of financial risk) briefly becomes very important at the height of the financial crisis in the autumn of 2008. Figure 8 shows that the CER future variable is increasingly an important price driver in Phase 2, while Figure 7 shows that the US carbon market variable (pcarbus) is a moderately strong price driver over the whole data span. Figure 9, which plots the error variance, shows a large increase in volatility in two time periods we have discussed above. These are the periods that correspond to the spring 2006 when the initial emissions verification process revealed the over-allocation of permits; and early 2009, when fraudulent activity affected the market (Frunza et al. 2010).

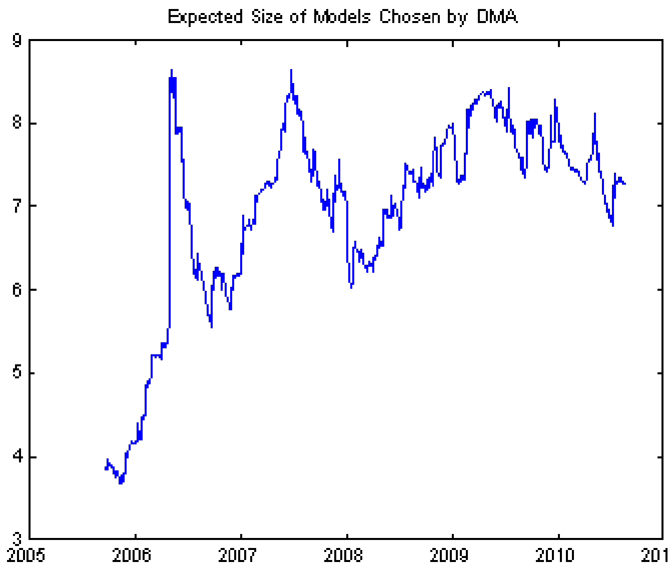

Figure 4: Results for Carbon Futures Price

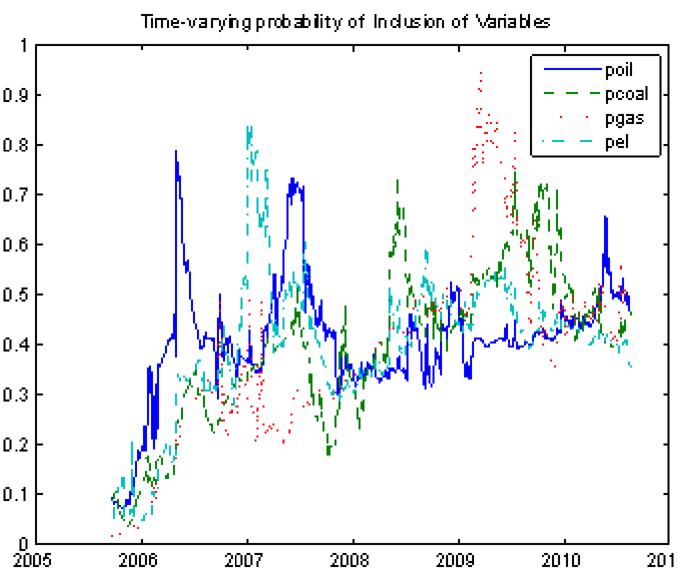

Figure 5: Results for Carbon Futures Price

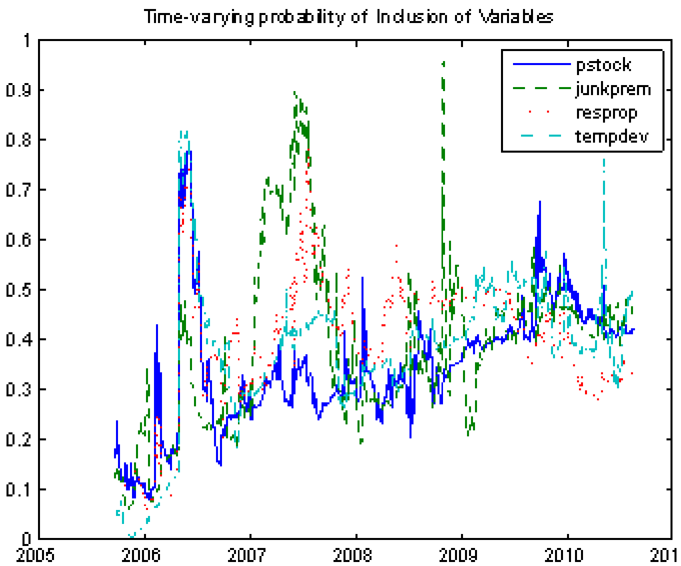

Figure 6: Results for Carbon Futures Price

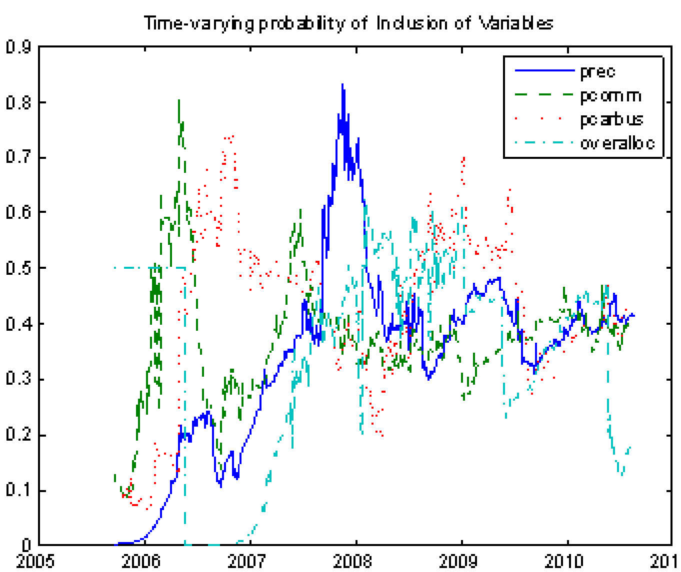

Figure 7: Results for Carbon Futures Price

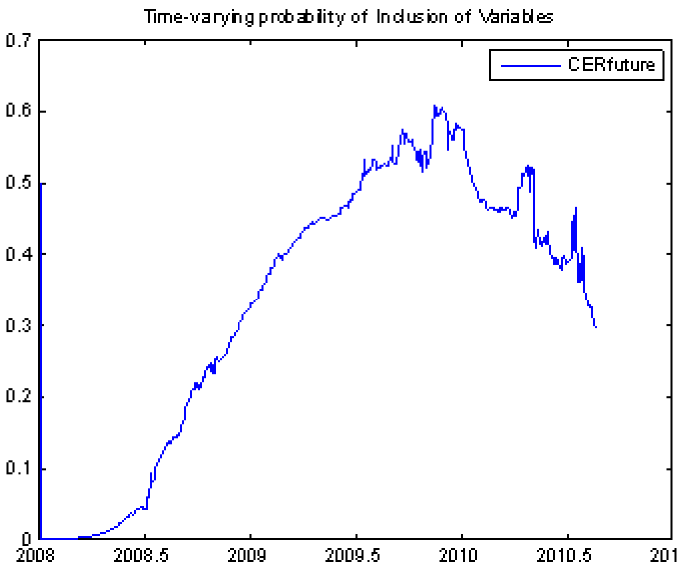

Figure 8: Results for Carbon Futures Price

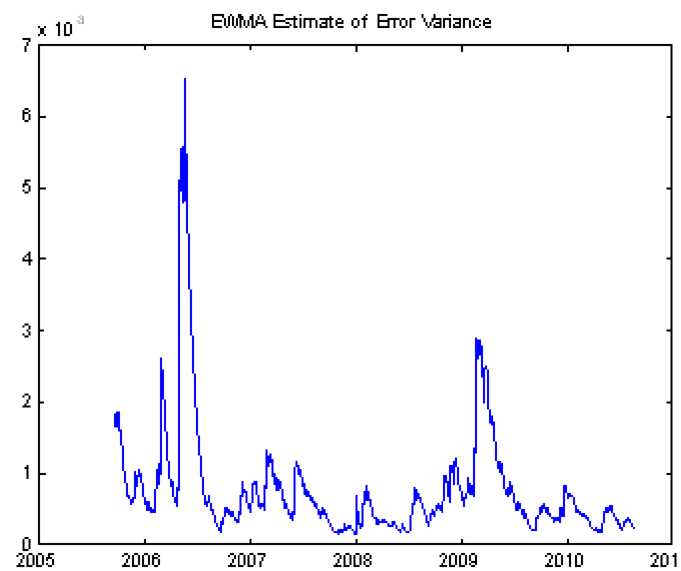

Figure 9: Results for Carbon Futures Price

Figures 10 through 13 present the same information as Figures 4 through 7, except that they use the log of the spot price instead of the log of the futures price as the dependent variable. Broadly speaking, we find a similar pattern of change in respect to price drivers and expected model size. However, there are some differences that are worth noting. First of all, we have not presented results for the CER future price driver analogous to Figure 8. The figure is basically a horizontal line at one. Thus, when the spot price is the dependent variable, DMA always wants to include the lag of the CER future price as a predictor. Remember that a lag of the spot price is also included in the model, so this result says that the lag of the CER future is providing additional explanatory power. Secondly, DMA is now choosing slightly less parsimonious models. A careful examination of Figures 10 through 13 indicates that variables such as the price of gas and tempdev (our measure of temperature extremes) have a more important role in respect to the spot than the futures price. This is to be expected since the spot price of a carbon permit is expected to reflect short term issues to a larger extent than a futures price with settlement date in 2012. For instance, an unusually hot summer will increase demand for carbon permits and thus affect the spot price. However an unusually hot summer will have less impact on the futures price. Similarly, the price of gas is an important factor bearing on the clean spark spread and, thus, on a firm's decision whether to abate or buy carbon permits to cover emissions. Changes in the clean spark price on a particular day will more likely impact the spot price than the futures price.

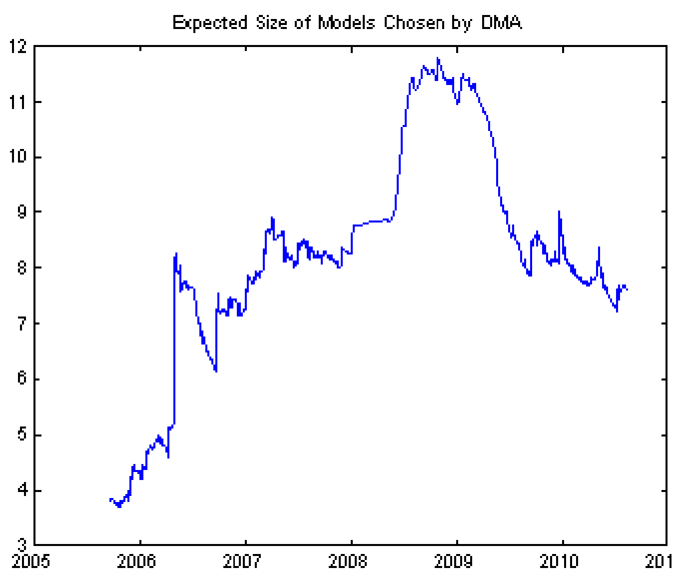

Figure 10: Results for Carbon Spot Price Data

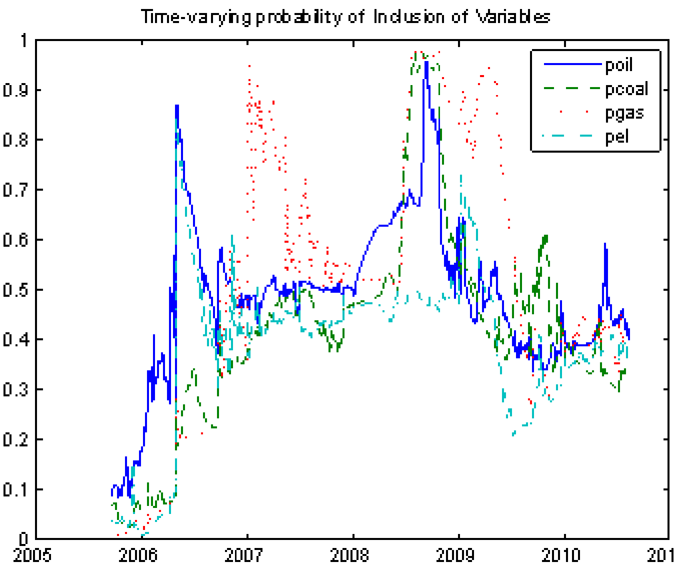

Figure 11: Results for Carbon Spot Price Data

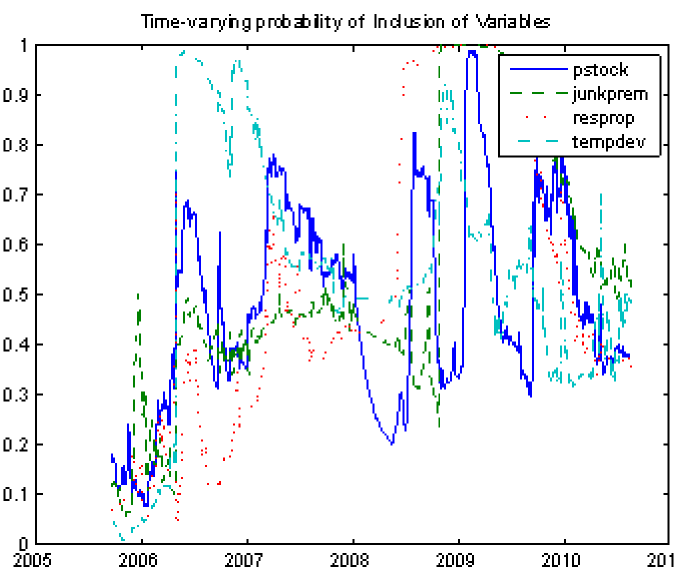

Figure 12: Results for Carbon Spot Price Data

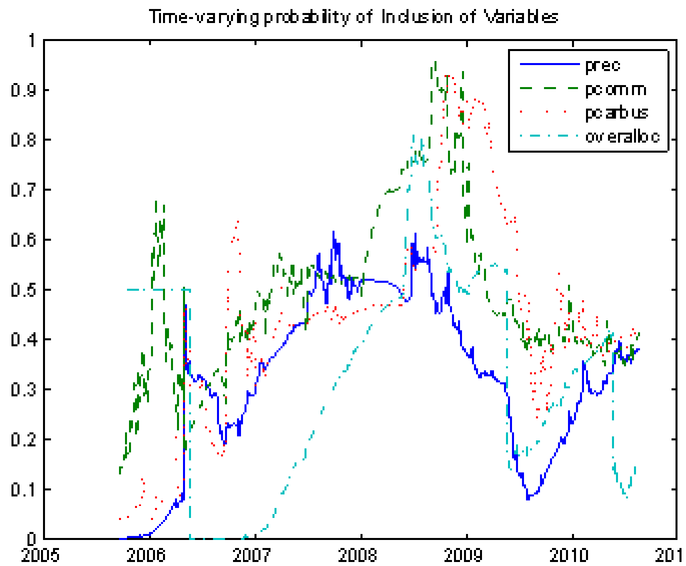

Figure 13: Results for Carbon Spot Price Data

5.3 Results of Forecasting Exercise

To evaluate the forecasting performance of DMA in the carbon market, we use mean squared forecast errors (MSFEs) and sums of log predictive likelihoods. The MSFE is a standard measure of the performance of point forecasts, whereas the sum of log predictive likelihoods evaluates the forecast performance of the entire predictive density. Formally, the predictive likelihood at time t is the predictive density (given information through $t-1$) evaluated at the actual outcome. DMA provides us with the predictive density (see Appendix B). For both MSFEs and sums of log predictive likelihoods we set aside an initial 100 daily observations and evaluate forecast performance on the remaining observations.

We compare our DMA forecasts with three other forecasting procedures. The first of these performs model averaging in the same manner as DMA, but is not a TVP model. That is, in (1) it sets $Q_t = 0$ so that coefficients do not change over time. The second is a TVP regression model including all the price drivers. No model averaging is involved, but rather this single model is used. The third is simply the single model which includes all the price drivers and does not allow for time variation in coefficients (i.e. it is analogous to a recursive OLS forecasting strategy using a regression model containing all the explanatory variables). Full details of all approaches are provided in Appendix B.

Table 1 compares the forecast performance of these different approaches.

Since low values of MSFEs and high values of predictive likelihoods, respectively, indicate good forecast performance, it can be seen that constant coefficient models always perform poorly. For the spot price data, which exhibits the pathologies seen in Figure 1, this is unsurprising. However, such a finding is less expected in the futures market. Overall we find DMA to consistently forecast well. However, in the spot market the TVP model also forecasts well (at least in terms of MSFE) whereas in the futures market DMA without TVP also performs well. DMA combines two aspects: it allows for parameters to change and for models to change. In the spot market, the first of seems very important; whereas in the futures market, the second is. Since in any empirical financial forecasting exercise either aspect could be important, the benefits of using DMA, which allows for both, are obvious.

Table 1: Forecasting Carbon Prices

	Spot		Futures	
	MSFE	Pred. Like.	MSFE	Pred. Like
DMA	116.15	1707.87	0.95	2822.81
DMA (no TVP)	233.73	426.84	0.94	2838.36
TVP (no DMA)	114.78	1664.12	1.11	2754.59
Const. Coeff.	250.39	358.06	1.02	2792.73

To understand how the results in Table 1 arise, it is worth studying the forecast performance over time in more detail. An important issue is how DMA handled the abrupt break in the spot price when we switched from Phase 1 to Phase 2. When forecasting the spot price on 1 January, 2008, the real-time forecaster only has information available through 31 December, 2007. The inclusion of a dummy variable for Phase 2 will be of no immediate benefit since on 31 December, 2007 there will be no data available to estimate its coefficient. On 31 December, 2007 the log of the spot price was -3.91, before switching to 3.18, 3.20, 3.20 and 3.21 on the first four trading days of January 2008. The DMA point forecasts on these five days are -3.94, -4.11, 10.24, 3.27 and 3.20. In other words, DMA (unsurprisingly) forecasts very poorly on the first two trading days of January, but by the third day was forecasting reasonably well and by the fourth and beyond was exhibiting very good forecast performance. We find this strong evidence that DMA can effectively handle even large structural breaks in financial time series. That is, it only took DMA three or four days to adjust its forecast performance to a break of huge magnitude in the dependent variable.

Table 1 contains sums of log predictive likelihoods over the entire sample. Figures 14 and 15 plot cumulative sums up to each point in time labelled on the X-axis for the futures and spot prices, respectively. A straight line in such a graph indicates that forecast performance is of a constant quality throughout the sample. For the futures data, apart from a small deterioration in forecast performance in May 2006 (i.e. the time when prices declined after the initial emissions verification revealed an over-allocation of permits), Figure 14 reveals such a straight line. For the spot price data, a clear deterioration in forecast performance is observed both in May 2006 and in early 2008 after the switch between Phase 1 and Phase 2.

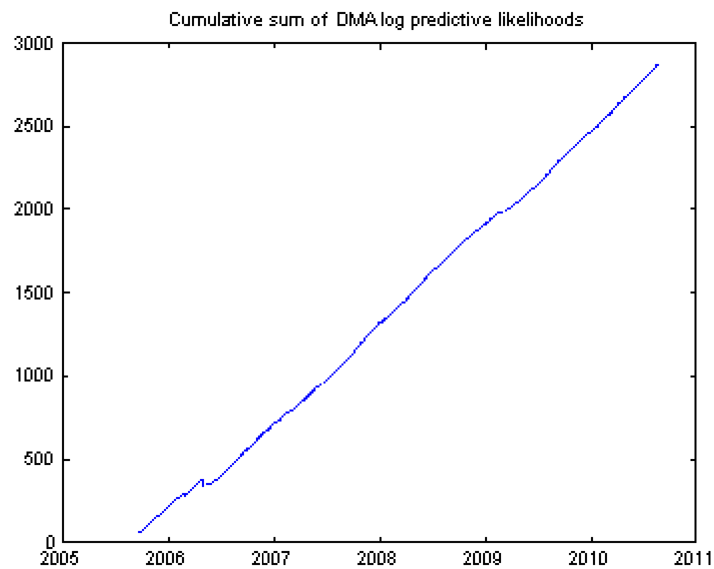

Figure 14: Results for Carbon Futures Price

5.4 Results Using only Phase 2 Data

Above, we have presented results using daily data from 22 April, 2005 through 18 August, 2010, which spans Phase 1 and Phase 2 of the EU ETS. We have argued previously that, even for a financial market hat has exhibited changes in its structure, it is best to include all the available data and use a method such as DMA to model the changes in structure. However, many (but not all) the papers in the literature use data from only one phase. Hence as a robustness check, we present results using exactly the same methods as above, but only using Phase 2 data. Table 2 presents results in the same format as Table 1. For the spot prices, DMA exhibits clearly superior forecast performance to all other methods. Thus, the superior forecast performance of DMA found in Table 1 was not due solely to its capability to handle the large break in spot prices that occurred in the switch between Phase 1 and Phase 2. For futures prices, the pattern of results is the same as in Table 1. DMA forecasts well, but the TVP aspect of the model adds little (since DMA without TVP also forecasts well). However, constant coefficient models forecast poorly. As with Table 1, we are finding strong evidence of the necessity of DMA (or a similar method which allows for parameter and model change) in order to forecast well in the carbon markets.

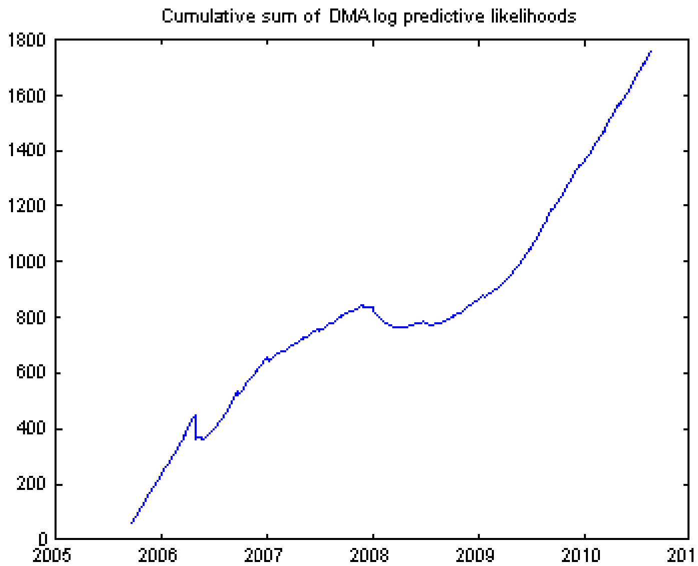

Figure 1: Figure 15: Results for Carbon Spot Price Data

Table 2: Forecasting Phase 2 Carbon Prices				
	Spot		Futures	
	MSFE	Pred. Like.	MSFE	Pred. Like
DMA	0.44	1268.79	0.38	1309.59
DMA (no TVP)	0.55	1229.12	0.39	1312.54
TVP (no DMA)	0.65	1066.27	0.42	1195.88
Const. Coeff.	0.99	1010.36	0.42	1217.06

For the sake of brevity, we will not produce figures comparable to Figure 4 through 15 based on Phase 2 data. Suffice it to note here that the general patterns are similar to those found using the entire sample: DMA does not provide us with a simple story where one or two price drivers are playing a predominant role. Rather the importance of the various price drivers is changing over time. Most of the specific patterns found using the entire sample are retained in the Phase 2 sample. For instance, the gas price remains an important price driver, particularly in the first half of 2009. At the time of the credit crunch, in the autumn of 2008, the junk bond premium becomes an important predictor. With regards to the error volatility, its increase in the first half of 2009 noted previously is replicated in the Phase 2 data.

5.5 Further Discussion of Results

The DMA methodology employed in this study has revealed a number of characteristics pertinent to understanding the properties of the carbon futures and spot markets. First, it does not seem to be the case that this market is operating efficiently. By way of explanation, note that the efficient markets hypothesis states that the price of an asset should reflect all available information about market fundamentals. Price changes should only occur in response to new information about the asset's fundamental value. In our study, the price drivers should measure new information. We are finding that the price drivers often do have forecasting power. However, in stable, mature and efficient markets, we would expect the role of the price drivers to be roughly constant over time. For example, in mature markets we might expect that every time the coal price rises it will impact the carbon permit price in roughly the same way. We are not finding this stability in our results.

Furthermore, results for the spot and future prices exhibit similar levels of instability. This is not what we would expect in mature markets. That is, financial theory suggests that futures (with settlement date December 2012) should be less sensitive to news that has a short run impact on the demand for carbon permits. However, we find, for example, that some of price drivers reflecting weather conditions (which would be expected to have most impact on the current year's demand for carbon permits) are important for the futures price even in Phase 1. This instability in the EU ETS could have impacts on emitters, making it more difficult for them to plan ahead and achieve reductions in CO_2 emissions in a cost effective manner (e.g. through investments in low-carbon technologies and optimal design of policies programs aimed at emissions reduction). Nonetheless, we do see how a method like DMA can adapt the forecasting model quite quickly over time to handle the instabilities in the carbon market. Tables 1 and 2 and Figures 14 and 15 show DMA to forecast well, with only brief deteriorations in forecast performance at the most turbulent times.

Second, while we hesitate to make any definitive statements about price drivers in a reduced form forecasting model with correlated explanatory variables, we also find that the importance of certain drivers switches over time: For instance, as in previous studies we find prices of gas, oil coal and electricity are important drivers that dominate the list of our variables, but that their predictive power is not sustained over the course of the study. We also find that some variables (e.g. the junk bond premium) gain particular importance suddenly, albeit only briefly, at the height of the financial crisis. The interpretation of the importance of such a variable as a price driver is not as straightforward as with conventional variables, such as temperature or energy. When global economic conditions became unstable, investment risks will increase (as measured by the premium demanded on risky bonds).Reflecting this risk, we find that the latter in turn impacted on the price of carbon. We also find the influence of carbon offsets on price to be very important. For the futures price, though, it is interesting to note that its influence began to decline near the end of our sample. This decline may have been due to a drop in issuance rates of CERs. Remember that the quantities of CERs requested by project developers from large scale, high-yield CER projects began to fall when they came under increasing scrutiny by the CDM.

The US voluntary carbon market variable (pcarbus) is a moderately strong price driver over the whole data span of the study. On one level, this is to be expected. Given the demand for emissions rights that US involvement in a global mandatory cap-and-trade scheme would generate, the strength of the voluntary market could be rightly taken to be a signal of the potential commitment of US businesses to reduce carbon emissions. Its abrupt fall in influence as a price driver for carbon futures in June 2009 may have reflected the fact that the price of a permit on the voluntary market plunged well below 60 cents. However, despite prices falling to as low as 10 cents after this period, it regained some of its status as a price driver, perhaps due to optimism over a long-awaited climate bill that failed to pass in July 2010.

Our finding that price drivers are switching over time suggests a market that is still unstable and immature. This failure may be as much due to the scheme's institutional features and related events as it is to market fundamentals. As the end of Phase 2 draws closer, a number of institutional issues continue to create problems for the market. These include: the uncertainty over the use and number of CER/JI offsets and emission reduction goals, on-going fraud in the marketplace, and the potential for carbon leakage of energy-intensive industries. Regarding the latter, there is worry that some carbon intensive businesses will not be able to pass on the cost of carbon to consumers or obtain sufficient permits at home due to a restriction in offsets and thus will move abroad.

All of these issues are likely to affect the efficiency of the carbon market in a number of ways: Uncertainty over caps and offsets leads to price instability and depresses investment in abatement and low carbon technologies. Moreover, the restriction of international offsets well below a supplementarity limit specified by the Kyoto Protocol means that reductions will not be done in a cost-effective way since the differences in marginal abatement costs between countries may no longer be exploited. As well as leading to declining investments in developing and emerging market countries, the restriction of offsets and the lack of an internationally binding global emissions reduction agreement means that it is unlikely that a single, global price for carbon will be achieved in the near future. This failure will continue to depress the EU ETS market price. Similarly, carbon leakage will be exacerbated by the restriction in use of international offsets, especially in energy-intensive sectors, with negative consequences for carbon markets. Finally, the EU's mandatory target of a 20% share for renewable energies by 2020 will be harder to achieve in the power industry, with restrictions on such offsets and no free allocations after 2012. Restrictions would be expected to lead to a rise in the price of carbon price. Conversely, if the EC subsidizes and provides other incentives for renewables, then the carbon price may actually fall as a large amount of emission reductions will be achieved outside of the market.

6 Conclusion

This paper has used a technique called DMA to forecast spot and futures prices in the EU ETS carbon market. Forecasting prices as accurately as possible is necessary for installations subject to compliance to make appropriate investment decisions based on long-term price expectations and for policymakers to design appropriate emission reductions policies and calculate the real costs of emissions reduction to society. We have argued that DMA is ideally suited for studying such a market, since it deals with problems caused by the proliferation of price drivers and allows for the changing effects of price drivers in the market and for the forecasting model itself to change over time. We find strong evidence of substantial turbulence and change in the EU ETS market. We show how our DMA approach can model these features and forecast accurately compared to other approaches. Finally, we find that the importance of price drivers in the EU ETS is constantly changing, although several retain their importance longer than do others.

References

Alberola, E., Chevallier, J., and Cheze, B. (2008a) The EU emissions trading scheme: disentangling the effects of industrial production and CO_2 emissions on carbon prices. *Int. Econ.* 116, 93-126.

Alberola, E., Chevallier, J., and Chèze, B. (2008b) Price drivers and structural breaks in European carbon prices 2005–2007. *Energy Pol.* 36, 787–797.

Alberola, E., Chevallier, J., and Chèze, B. (2009a) Emissions compliances and carbon prices under the EU ETS: A country specific analysis of industrial sectors. *J. Pol. Modeling,* 31, 446-462.

Alberola, E., Chevallier, J. (2009b) European carbon prices and banking restrictions: evidence from Phase 1 (2005-2007). *Energy J.* 30, 51-80.

Asche, F., Osmundsen, P. and Sandsmark, M. (2006) The UK market for natural gas, oil and electricity: are the prices decoupled?. *Energy J.* 27, 27-40

Avramov, D. (2002) Stock return predictability and model uncertainty. *J. Finan. Econ.* 64 423–458

BBC (2011) "Phishing attack nets 3 million euros of carbon permits". February 9th. Available through the website http://news.bbc.co.uk.

Böhringer, C. and Lange, A. (2005) On the design of optimal grandfathering schemes for emission allowances. *Europ. Econ. Rev.* 49, 2041-2055.

Bredin, D. and Muckley, C. (2010) An emerging equilibrium in the EU emissions trading scheme. *Energy Econ.* forthcoming.

Bunn, D. and Fezzi, C. (2007) Interaction of European carbon trading and energy prices, *Fondazione Eni Enrico Mattei Working Papers*, Paper 123.

Chevallier, J. (2009) Carbon futures and macroeconomic risk factors: A view from the EU ETS. *Energy Econ.* 31, 614-625.

Christiansen, A., Arvanitakis, A., Tangen, K., and Hasselknippe, H., (2005) Price determinants in the EU emissions trading scheme. *Climate Pol.* 5, 15-30.

Cogley T. and Sargent T. (2005) Drifts and volatilities: Monetary policies and outcomes in the post WWII U.S. *Review of Economic Dynamics*, 8, 262-302.

Cogley T, Morozov S. and Sargent T. (2005) Bayesian fan charts for U.K. inflation: Forecasting and sources of uncertainty in an evolving monetary system. *J. Econ. Dynam. Control* 29, 1893-1925.

Convery, F., Redmond, L. (2007) Market and price developments in the European Union Emissions Trading Scheme. *Rev. Environ. Econ. Policy* 1, 88–111.

Crampton, P. and Kerr, S. (2002) Tradable carbon permit auctions. How and why to auction not grandfather. *Energy Pol.* 30, 333-345.

Eicher, T., Papageorgiou, C. and Raftery, A. (2010) Determining growth determinants: Default priors and predictive performance in Bayesian model averaging, with applications to growth determinants. *J. Appl. Econometrics* , forthcoming.

Ellerman, A. and Buchner, B. (2008) Over-allocation or abatement? A preliminary analysis of the EU ETS based on the 2005-06 emissions data Environ. Resource Econ. 41, 267-287.

Fernandez, C., Ley, E. and Steel, M. (2001a) Model uncertainty in cross-country growth regressions. *J. Appl. Econometrics* 16, 563-576.

Fernandez, C., Ley, E. and Steel, M. (2001b) Benchmark priors for Bayesian mode averaging. *J. Econometrics* 100, 381-427.

Frunza, M.-C., Guegan, D. and Thiebaut, D. (2010) Missing trader fraud on the emissions market, *Documents de Travail du Centre d'Economie de la Sorbonne*, 2010.71.

Hintermann, B. (2010) Allowance price drivers in the first phase of the EU ETS. *J. Environ. Econ. Manage.* 59, 43-56.

IETA (International Emissions Trading Association) (2010). GHG Market Report 2010. Available through the website http://www.ieta.org.

Josse-Vasquez, C.I., and Neumann, A. (2006). Transatlantic natural gas and oil price relationships: an empirical analysis, manuscript, Dresden University of Technology

Kanen, J.L.M. (2006). *Carbon Trading and Pricing.* Environmental Finance Publications: London, UK.

Klein Tank, A.M.G., Wijngaard, J., Können, G., Böhm, R., Demarée, G., Gocheva, A., Mileta, M., Pashiardis, S., Hejkrlik, L., Kern-Hansen, C., Heino, R., Bessemoulin, P., Müller-Westermeier, G., Tzanakou, M., Szalai, S., Pálsdóttir, T., Fitzgerald, D., Rubin, S., Capaldo, M., Maugeri, M., Leitass, A., Bukantis, A., Aberfeld, R., van Engelen, A., Forland, E., Mietus, M., Coelho, F., Mares, C., Razuvaev, V., Nieplova, E., Cegnar, T., Antonio López, J., Dahlström, B., Moberg, A., Kirchhofer, W., Ceylan, A., Pachaliuk, O., Alexander, L. and Petrovic, P. (2002) Daily dataset of 20th-century surface air temperature and precipitation series for the European Climate Assessment. *Internat J. Climatology.* 22, 1441-1453. Data and metadata available at http://eca.knmi.nl.

Koop G and Korobilis, D. (2010) Forecasting inflation using dynamic model averaging. Manuscript available at
http://personal.strath.ac.uk/gary.koop/koop_korobilis_forecasting_inflation_using_DMA.pdf.

Kossoy, A. and Ambrosi, P. (2010) *State and trends of the carbon market 2010.* World Bank: Washington, D.C.

Ley, E. and Steel, M. (2009) On th eeffect of prior assumptions in Bayesian model averaging with applications to growth regression. *J. Appl. Econometrics* 24, 651-674.

Ley, E. and Steel, M. (2010) Mixtures of g-priors for Bayesian model averaging with economic applications, Centre for Research in Statistical Methology, Working Paper No. 10-23.

Mansanet-Bataller, M., Chevallier, J., Hervé-Mignucci, M. and Alberola, E. (2010) The EUA-sCER Spread: Compliance strategies and arbitrage in the European carbon market, Mission Climat Working Paper, 2010-6. Available through the website http://www.caissedesdepots.fr.

New York Times (2011) "Emission permits theft estimated at $37 million" January 20. Available through the website http://www.nytimes.com.

Primiceri G. (2005) Time varying structural vector autoregressions and monetary policy. *Rev. Econ. Stud.* 72, 821-852.

Raftery, A., Karny, M. and Ettler, P. (2010) Online prediction under model uncertainty via dynamic model averaging: Application to a cold rolling mill, *Technometrics*, 52, 52-66.

Raftery, A., Madigan, D. and Hoeting, J. (1997) Bayesian model averaging for linear regression models. *J. Amer.Stat. Assoc.* 92, 179-191.

RiskMetrics. (1996) Technical Document (Fourth Edition) available at http://www.riskmetrics.com/system/files/private/td4e.pdf.

Appendix A: Data

Independent Variables Summary			
Variable Description	*Acronym*	*Source*	*Measure*
Phase 1 & 2 Dummy	dumphase1	–	1=Phase1; 0=Phase 2
Emissions-to-cap	overalloc	EU	Total annual EU actual emissions minus total annual EU allocated permits
Oil price	poil	ICE	daily price brent oil euros/barrel
Coal price	pcoal	McCloskey Coal	NW Europe Coal Marker (average price at ARA hub)
Gas price	pgas	APX-ENDEX	TTF continental gas price
Electricity price	pel	EEX	Phelix Base base load price
Temperature	tempdev	European Climate Assessment	Absolute deviation from daily population weighted mean temperature
Hydropower capacity	resprop	Noordpool Ex	Nordic reservoir capacity
Precipitation	precip	European Climate Assessment	Nordic rainfall levels in mm
Stock Market Performance	pstock	Thompson's Datastream	Euronext-100 stock market index
Commodity Prices	prcomm	Thompson's Datastream	Reuters Commodities Index
Corporate Risk Premium	junkprem	Thompson's Datastream	Spread between Moody's BAA and AAA bonds
US carbon market	pcarbus	CCX	Price carbon permits, US$/tonne
Carbon offset price	CERfuture	ECX	Price CER permit, euros/tonne

We use data from 22 April, 2005 through 18 August, 2010.There are missing values in the spot price of a permit from 1 January, 2008 through 25 February, 2008. The cause of this was a delay in issuing the Phase 2 permits.[20] For these time periods, we use the price of a future permit with closest settlement date (i.e. yearly settlement dates). This procedure can be justified by noting that, for periods when data on both the spot price and the future with yearly settlement date are available, the correlation between them is 0.9995 and they are virtually the same as one another.

Also note the following about the variable measuring overallocation of permits: At each point in time, we take the most recent allocation and verification of emissions that would be known at that point in time and subtract the latter from the former. Verified emissions data for the previous year is released on 15 May. Thus, for example, on 14 May, 2007 the forecaster would only have the 2005 overallocation data, but on 15 May, 2007 it becomes the 2006 overallocation data.

[20] This delay occurred because the EU's CITL was not connected to the UN's International Transaction Log (ITL), which caused some governments to delay issuance of new permits. For more explanation, see http://www.carbonfinanceonline.com/index.cfm?section=lead&action=view&id=11081.

Appendix B: Statistical Methodology (DMA)

We work with the set of models defined in (2). Estimation within a single model involves familiar statistical methods involving the Kalman filter and will not be explained here in detail. Suffice it to note that (apart from our estimation of $H_t^{(k)}$ which will be discussed below), our estimation methods within a single model are exactly as in Raftery et al. (2010). This involves the choice of a forgetting factor, λ, and we (following Raftery et al., 2010) set $\lambda = 0.99$. The use of a forgetting factor implicitly estimates Q_t^k. With regards to initialization of the Kalman filter, we use a diffuse prior $\theta_0^{(k)} \sim N\left(0, 100 I_{m_k}\right)$, where m_k is the number of variables in model k, for $k = 1, .., K$.

To simplify notation, let $\pi_{t|s,k} = \Pr\left(L_t = k | y^s\right)$. The new recursions required by DMA involve (beginning with $\pi_{0|0,k}$) $\pi_{t|t-1,k}$ and $\pi_{t|t,k}$. DMA proceeds by averaging across forecasts using $\pi_{t|t-1,k}$ as weights for $k = 1, .., K$ and $t = 1, .., T$. To be precise, DMA point forecasts are given by:

$$E\left(y_t | y^{t-1}\right) = \sum_{k=1}^{K} \pi_{t|t-1,k} Z_{t-1}^{(k)} \widehat{\theta}_{t-1}^{(k)},$$

where $\widehat{\theta}_{t-1}^{(k)}$ is the Kalman filter estimate of the regression coefficients given information available at time $t-1$.

Raftery et al. (2010) justify and use:

$$\pi_{t|t-1,k} = \frac{\pi_{t-1|t-1,k}^{\alpha}}{\sum_{l=1}^{K} \pi_{t-1|t-1,l}^{\alpha}}, \tag{3}$$

where $0 < \alpha \leq 1$ is a forgetting factor which is set to a fixed value slightly less than one. The huge advantage of using the forgetting factor α in the model prediction equation is that we do not require an MCMC algorithm to draw transitions between models nor a simulation algorithm over model space. Instead, simple evaluations comparable to those of the updating equation in the Kalman filter can be done. In particular, we have a model updating equation of:

$$\pi_{t|t,k} = \frac{\pi_{t|t-1,k} p_k\left(y_t | y^{t-1}\right)}{\sum_{l=1}^{K} \pi_{t|t-1,l} p_l\left(y_t | y^{t-1}\right)}, \tag{4}$$

where $p_l\left(y_t | y^{t-1}\right)$ is the predictive density for model l (i.e. this is the standard Normal predictive density provided by the Kalman filter) evaluated at y_t.

To understand further how the forgetting factor α can be interpreted, note that this specification implies that the weight used in DMA which is attached to model k at time t is:

$$\begin{aligned} \pi_{t|t-1,k} &\propto \left[\pi_{t-1|t-2,k} p_k\left(y_{t-1} | y^{t-2}\right)\right]^{\alpha} \\ &= \prod_{i=1}^{t-1} \left[p_k\left(y_{t-i} | y^{t-i-1}\right)\right]^{\alpha^i}. \end{aligned}$$

Thus, model k will receive more weight at time t if it has forecast well in the recent past (where forecast performance is measured by the predictive density, $p_k\left(y_{t-i}|y^{t-i-1}\right)$). The interpretation of "recent past" is controlled by the forgetting factor, α and we have an exponential decay at the rate α^i for observations i periods ago. Thus, if $\alpha = 0.99$ (the value used by Raftery et al., 2010), forecast performance four weeks ago receives 80% as much weight as forecast performance last period (when using daily data based on a five day week).

In our forecasting exercise, we compare DMA to several alternatives. These are obtained as special cases of DMA. Our first comparator does model averaging, but does not allow for time variation in parameters. This is obtained by setting $\lambda = 1$, but all other assumptions are identical to those used with DMA. The second comparator allows for time variation in parameters, but does not do model averaging. Instead it simply works with a single model which includes all of the explanatory variable. This is obtained by restricting the DMA model space to this single model, but all other assumptions are as with DMA. Our third comparator allows for no model averaging nor time variation in parameters (beyond that associated with any recursive forecasting exercise). It combines the assumptions of our first and second comparators (i.e. it works with the same single model as the second, but sets $\lambda = 1$ as in the first).

We stress that, conditional on H_t, the estimation and forecasting strategy outlined above only involves evaluating formulae such as those in the Kalman filter. The recursions above are started by choosing $\pi_{0|0,k}$ for $k = 1,..,K$ and we make the noninformative choice of $\pi_{0|0,k} = \frac{1}{K}$ for $k = 1,..,K$.

The preceding discussion is all conditional on H_t. Raftery et al. (2010) recommend a simple plug in method where $H_t^{(k)} = H^{(k)}$ and is replaced with a consistent estimate. When forecasting financial variables, however, it is likely that the error variance is changing over time. Thus, we use an Exponentially Weighted Moving Average (EWMA) estimate of $H_t^{(k)}$:

$$\widehat{H}_t^{(k)} = \sqrt{(1-\kappa)\sum_{j=1}^{t}\kappa^{j-1}\left(y_j - z_j^{(k)}\widehat{\theta}_j^{(k)}\right)^2}. \tag{5}$$

EWMA estimators are commonly used to model time-varying volatilities in finance; see Riskmetrics (1996) for the properties of EWMA estimators. κ is called a decay factor, and Riskmetrics proposes setting 0.94 for daily data and we follow this choice. An attractive feature of the EWMA specification is that it can be approximated by a recursive form, which can be used to obtain volatility forecasts. The period $t+1$ forecast given data up to time t takes the form.

$$\widehat{H}_{t+1|t}^{(k)} = \kappa\widehat{H}_{t|t-1}^{(k)} + (1-\kappa)\left(y_t - Z_t^{(k)}\widehat{\theta}_t^{(k)}\right)^2.$$