

**Rules vs. Discretion in Cap-
and-Trade Programs:
Evidence from the EU
Emission Trading System**

Marina Friedrich, Sébastien Fries, Michael Pahle, Ottmar Edenhofer

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Rules vs. Discretion in Cap-and-Trade Programs: Evidence from the EU Emission Trading System

Abstract

Long-term commitment is crucial for the dynamic efficiency of intertemporal cap-and-trade programs. Discretionary interventions in such programs could destabilize the market, and necessitate subsequent corrective interventions that instigate regulatory instability (Kydland and Prescott, 1977). In this work, we provide evidence for this claim from the EU's cap-and-trade program (EU-ETS). We ground our analysis in the theoretical finance literature, and apply a mixed method approach (time-varying regression, bubble detection, crash-odds modelling). We find that the recent EU-ETS reform triggered market participants into speculation, which likely led to an overreaction that destabilized the market. We discuss how the smokescreen politics behind the reform, which manifested itself in complex rules, was crucial for this outcome. We conclude that rules only ensure long-term commitment when their impact on prices is predictable.

JEL-Codes: Q480, Q500, Q560.

Keywords: rules vs. discretion, cap-and-trade, overreaction, allowance pricing.

*Marina Friedrich**

*Vrije Universiteit Amsterdam
Department of Econometrics & Data Science
De Boelelaan 1105
The Netherlands - 1081HV Amsterdam
m.friedrich@vu.nl*

Sébastien Fries

*Vrije Universiteit Amsterdam
Department of Econometrics & Data
Science, De Boelelaan 1105
The Netherlands - 1081HV Amsterdam
s.f.fries@vu.nl*

Michael Pahle

*Potsdam Institute for Climate Impact Research
Member of the Leibniz Association
Potsdam / Germany
michael.pahle@pik-potsdam.de*

Ottmar Edenhofer

*Potsdam Institute for Climate Impact
Research, Member of the Leibniz
Association, Potsdam / Germany
ottmar.edenhofer@pik-potsdam.de*

*corresponding author

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

A previous version of this paper entitled Allowance prices in the EU ETS – fundamental price drivers and the recent upward trend has been presented at the conferences Econometric Models of Climate Change III in Rome and EAERE in Manchester. The authors thank the participants for helpful discussions. The authors additionally thank Samuel Okullo, Oliver Tietjen, Sean Telg, Mikkel Bennedsen, Simon Quemin, Dallas Burtraw, and Linus Mattauch for valuable comments. The research leading to these results has received funding from BMBF under the research project START and the European Union's Horizon 2020 research and innovation programme under grant agreement No 730403 (INNOPATHS).

1 Introduction

The dynamic efficiency of intertemporal cap-and-trade programs stands and falls with long-term commitment. If credible, it can solve the problem of distorting investment decisions through the ratchet effect, the hold-up problem, and the problem of inappropriate risk allocation (Hepburn, 2006). Yet, regulators want to be flexible to cope with future circumstances, which is why they prefer discretionary intervention to commitment (Hasegawa and Salant, 2014). This can result in economic instability: Firms may persistently mispredict the intervention's long-term impact on allowance prices, which may necessitate further correcting interventions in the future (Kydland and Prescott, 1977). The anticipation of this may create regulatory uncertainty which, in turn, further distorts allowance prices (Salant and Henderson, 1978; Salant, 2016). Contrary to their intention, discretionary interventions in cap-and-trade programs could thus have the perverse effect of destabilizing the allowance market.

While the theoretical case is clear, empirical evidence is, to the best of our knowledge, still missing. Several papers have analyzed the effects of regulatory announcements in the EU Emissions Trading System (EU ETS) using event studies. They find significant effects on the prices of European Emission Allowances (EUAs) for many of the considered announcements (Koch et al., 2016; Deeney et al., 2016). However, due to the nature of the method, such studies focus on relatively narrow time windows and nothing can be said about enduring effects on the market. Moreover, all analyzed announcements were related to relatively minor, cap-neutral reforms, i.e. they did not affect the total number of allowances. The most recent reform, enacted in March 2018, revised the market rules for the coming decade and




Figure 1: EUA price development – nearest December Futures Contract
Source: EEX Leipzig; see Section 3.1 for details.

entails an adjustment of the cap. Around the same time, an unprecedented price run up occurred as displayed in Figure 1. This price increase is widely attributed to the reform and it provides us with a unique opportunity to empirically analyze the impact of a discretionary intervention.

Our work expands the literature in two key ways. First, we study market interventions from a different angle. Contrary to previous work, our focus is on the impact of long-term market stability rather than short-term price responses. Acknowledging that allowance markets are increasingly recognized as financial markets (Friedrich et al., 2020), we ground our analysis in the theoretical finance literature. This literature emphasizes the important role of speculation and belief formation for market stability. From it we derive the hypothesis that discretionary interventions constitute (good) news about fundamentals that could lead to an overreaction of the market indicated by explosive prices - similar to what is observed during the inflationary period of bubbles (Diba and Grossman, 1988; Abreu and Brunnermeier, 2003; Barberis et al., 2018). Second and related, we employ

a novel mixed-method approach to rigorously investigate the overreaction hypothesis.

In our first step, we use a flexible time-varying regression approach to investigate whether the price run-up is accompanied by a parallel development in the role of market fundamentals. In a second step, we formally test for periods of explosive behaviour in allowance prices with the help of an econometric testing procedure which was proposed by Phillips et al. (2015). The date stamping approach that is part of this procedure allows us to examine whether the adoption of the reform coincides with the onset of the explosive period. In a third step, we predict the crash odds of the explosive period found in the second step using the novel method of Fries (2018a). We compare the estimated crash odds to the ex-post evolution of prices.

We obtain three main results. First, during the price run-up period demand-side fundamentals have neither gained importance as price determinants nor do they show the same explosive behaviour as allowance prices. Second, the onset of the explosive period is time-stamped to March 2018, which is the month in which the reform was adopted. Third, the estimated crash odds of the explosive period after October 2018 are quite high, suggesting a likely price collapse within a year. All three results suggest that the reform triggered market participants into speculation about its price impacts, which likely led to an overreaction that destabilized the market. We further discuss out-of-sample price development in light of our findings, and how the smokescreen politics behind the reform was crucial for this outcome. We conclude that when implementing or adjusting rules in cap-and-trade programs, it is crucial to consider the beliefs they instill about

their price impacts. More specifically, our results imply that the better their effect on prices can be predicted, the lower the risk that they destabilize the market.

The paper is structured as follows. In Section 2, we provide the policy background and review the related literature. In Section 3, we introduce our dataset and describe the empirical strategy. Section 4 analyses the relationship between allowance prices and its fundamental price drivers. Subsequently, Section 5 investigates potential explosive behaviour in allowance prices. First, we test for explosive behaviour. Second, based on the findings which deliver evidence of such behaviour in our series, we model the explosive period and estimate probabilities of collapse at different time horizons. Section 6 concludes and discusses the implications.

2 Policy background and related literature

2.1 The recent 2018 EU-ETS reform

Persistently low prices from 2013 on led EU policy makers to initiate a series of reforms primarily intended, in their own words, to "reduce the surplus of emission allowances in the carbon market".¹ The reform process was concluded with an amendment of the EU ETS directive in March 2018. In essence, the reform entailed the adoption of two new market rules: (1) From 2019 on, allowances will be transferred to and from the so called Market Stability Reserve (MSR), depending on the total number of allowances in circulation in the market and a given intake/outtake threshold. (2) From 2024 on, a certain number of allowances in the MSR will be cancelled,

¹https://ec.europa.eu/clima/policies/ets/revision_en

depending on the number of allowances auctioned in the previous year.

These rules both retroactively changed the functioning of the EU ETS, and added considerable complexity to the market (Perino, 2018). In fact, this particular design was chosen deliberately in an act of "smokescreen politics" (Wetttestad and Jevnaker, 2019). By making the distribution of costs (from future cancellation) obscure and diffuse, policy makers could muster sufficient political support to pass the reform. The main drawback of this political strategy is that market participants can only speculate about the the price impacts of the new rules.

Surveys and anecdotal evidence suggest that speculations has indeed been an important price driver. Two recent surveys among market participants find that the anticipation of a more stringent market balance through the MSR and the higher Linear Reduction Factor are seen as important causes (ZEW, 2019; Thomson Reuters, 2019). However, in both surveys also speculative buying ranks high as a driver. Furthermore, media and market intelligence spread the information that much of the price run-up was due to trading activities from speculative investors, like banks and hedge funds, in anticipation of the (price) effect of the reform.² Indeed, overall trading volumes in 2018 increased by 42% compared to the previous year (Marcu et al., 2019). However, this does not rule out that market fundamentals have been an equally important price driver.

In the following, we review the empirical literature on fundamental price drivers and the literature on speculative bubbles. We briefly treat bubble formation and mention related papers using bubble testing as well as bubble modelling.

²<https://www.ft.com/content/6e60b6ec-b10b-11e8-99ca-68cf89602132on08.09.2018>.

2.2 Literature on fundamental and finance related price drivers

According to environmental economics, market fundamentals such as coal and gas prices as well as economic activity should have a major effect on allowance prices. A study of the related literature shows, however, that empirical evidence is mixed. Previous studies indicate that standard approaches, such as linear regression models, need to be adapted by splitting the sample, including breaks or dummy variables in order to improve their findings. This is discussed in a review of the empirical literature of price formation in the EU ETS by Friedrich et al. (2020).

Time constant models have frequently been used in this context; for instance, in Hintermann (2010), Koch et al. (2014) and Aatola et al. (2013). The question whether it might be more appropriate to account for potential time-variation when modelling the relationship between allowance prices and fundamentals has been raised in Lutz et al. (2013) who consider potential non-linearities. Using a regime-switching model, they distinguish two different pricing regimes - one applies during periods of high volatility and the other during periods of low volatility. By construction, the impact of explanatory variables on the allowance price can differ among the two regimes. In both regimes, they find the same set of relevant price drivers. Coal and gas prices, oil prices and the stock index are statistically significant determinants of the EUA price. In Regime 2, which is characterised by low and constant volatility, all significant price drivers show the anticipated sign. Regime 1, however, shows a positive impact of the coal price. This goes against economic considerations that predict, as in the second regime, a negative effect of the coal price on allowance prices. In a recent contribu-

tion, Jiménez-Rodríguez (2019) apply a time-varying coefficient model to assess the impact of stock price indices on the allowance price and find a time-varying effect. This paper differs from our work as it does not include other price drivers apart from the stock price indices. Moreover, a different model and estimation method are used and no confidence intervals are provided. Both papers, Lutz et al. (2013) and Jiménez-Rodríguez (2019), provide further evidence that the relationship between the allowance price and its fundamentals might not be constant over time but can be subject to (structural) changes.

The effect of the coal price on allowance prices causes further disagreement in findings. Similar to the results in the second regime of Lutz et al. (2013), Rickels et al. (2014) find a positive effect of the coal price on the allowance price in their single variable analysis. The paper by Rickels et al. (2014) differentiates itself from previous studies, because the authors investigate the effect of the choice of data series by performing various regressions with only one explanatory variable. They consider multiple data series for the different factors, from different sources and with different sampling frequency (daily and weekly). In their final regression specification, they do not include the coal price as a separate explanatory factor but as part of the switching price. Further evidence is found in Aatola et al. (2013) for the period 2005-2010 who find a negative coefficient of coal, while Hintermann (2010) finds it to be insignificant in Phase I data. In addition, Koch et al. (2014) look at the entire second Phase II and the first year of Phase III and find an insignificant coefficient of coal. However, the explicitly calculated fuel switching price is found to have a significant effect. Regarding the effect of the gas price on allowance prices, there is no ambiguity. All

studies find a positive and significant coefficient of the gas price independent which approach is used. In particular, in Hintermann (2010) it is the only explanatory variable with a significant effect throughout all considered specifications.

The finance literature takes a different angle on price formation and is in particular interested in situations where prices deviate from fundamentals. Of particular relevance here is the finding that good news – such as the ETS reform in our case – may trigger a bubble. The underlying mechanism could be, for example, price extrapolation (Barberis et al., 2018) and speculation (Abreu and Brunnermeier, 2003). A primary indicator for the inflationary period of a bubble is explosive behavior in prices, and recently a new method to detect and timestamp such periods was proposed by Phillips et al. (2015). This method has subsequently been applied by Corbet et al. (2018) to Bitcoin prices, by Shi (2017) to the US housing market, and by Sharma and Escobari (2018) to the energy sector.

The related paper by Creti and Joëts (2017) uses this method to investigate EU ETS prices. The authors find evidence of short explosive episodes in their sample from 2005 to 2014. The detected periods do not last longer than a few days – only 2 out of 11 last longer than 5 days, among them one downward period lasting 12 days and one upward period lasting 9 days. This finding indirectly supports our presumption that no reform (before 2014 in this case) had an impact on prices similar to the magnitude and persistence we suspect for the 2018 reform.

3 Data and methods

In this section, we first present our data set and carefully investigate the stationarity properties of each time series. This is an important starting point before applying our proposed methods as some of them require the data to be stationary. In the final part of this section, we describe our empirical strategy.

3.1 EUA prices and explanatory variables

We consider weekly data for the period from January 2008 to October 2018 resulting in $T = 538$ observations. This sample period covers the entire Phase II and a large part of Phase III. As allowance price series, we use the December futures contract traded on the European Energy Exchange (EEX). The contract is rolled over to the next contract at the end of October each year. The resulting continuous price series is displayed in Figure 1. Our sample captures the entire period of rapid increase in allowance prices which started in mid 2017. It ends with the first point of considerable decline at the end of October 2018. This allows us to investigate the entire trending period while leaving enough room for out-of-sample prediction of crash odds in the final step of our empirical approach.

We choose December futures prices like most related papers since they are frequently traded (e.g. Koch et al., 2014; Lutz et al., 2013; Aatola et al., 2013). As our main set of explanatory variables, we include natural gas and coal prices as month-ahead futures from the same platform as well as the stock index STOXX Europe 50 as an indicator of current and expected economic activity.³ As an alternative, we use data on a comparable index,

³Other commonly used indicators such as the Industrial Production Index or the

which is sometimes used in this context, the STOXX Europe 600 index.⁴ Further, we consider the oil price. In the related literature, there is no clear agreement on whether its effect is due to being a proxy for economic activity or if it comes from the (limited) fuel switching from oil to gas (Hintermann et al., 2016). In addition, we use daily mean temperature data for seven European cities from the European Climate Assessment & Dataset (ECA&D) presented in Klein Tank et al. (2002). We transform the data into a series of weekly temperature averages and we take out the seasonal component with the help of Fourier terms. Details on the location of the stations and the Fourier term regression can be found in Appendix B.4.


The gas price is the settlement price of month-ahead Dutch TTF futures, denoted in EUR/MWh. TTF stand for Title Transfer Facility and is a virtual trading point for natural gas in the Netherlands. The TTF futures contract is a frequently used series for gas prices in the related literature. Similarly, the coal price we consider is the settlement price of month-ahead futures based on the API2 index of the ARA region (Amsterdam-Rotterdam-Antwerp). The contract size is 1,000 metric tonnes of thermal coal.⁵ Both are obtained from the EEX. For oil we rely on the historical futures prices (continuous contract) of Brent crude oil based on data from the Intercontinental Exchange (ICE), retrieved from Quandl.⁶ The contract

Economic Sentiment Indicator are measured at a lower frequency such that we do not consider them here. Another promising indicator could be nighttime light from satellite observations.


⁴The stock index data are retrieved on 21.01.2019 from <https://quotes.wsj.com/index/XX/SXXP/historical-prices> (STOXX Eur 600) and <http://quotes.wsj.com/index/XX/SX5E/historical-prices> (STOXX Eur 50).

⁵To convert the coal price data into EUR/MWh, one simply has to divide the series by the conversion factor of 8.14. Since the conversion factor is constant and we consider first differences, this would not change our results.


⁶Retrieved from https://www.quandl.com/data/CHRIS/ICE_B1 on 21.01.2019




(a) Coal




(b) Gas




(c) Oil



(d) Stoxx 50



(e) Stoxx 600



(f) Temperature averages

Figure 2: Plot of the time series in our data set

Source: For a detailed description of the sources we refer to Section 3.1.

Notes: In Panels (a) to (c), prices are given in EUR, while the coal and the oil price was transformed from USD into EUR before plotting. Panels (d) and (e) are given in index values and (f) in degrees Celsius.

size is 1,000 barrels. The coal and the oil prices need to be converted into EUR, as they are denoted in USD. This is done using USD/EUR exchange rate data from Tullett Prebon.⁷

⁷Retrieved from <https://quotes.wsj.com/fx/EURUSD/historical-prices> on 21.01.2019

Various other variables – e.g. data on renewable energy production, issued Certified Emission Reductions and fuel switching prices – appear in the literature. Apart from the switching price, these variables are often found to be insignificant or to have a negligible effect in terms of the magnitude of the coefficients. Further limitations arise, since data on wind, solar or hydro production is only available for specific countries or regions. In our preliminary regression analysis we consider data on energy supply from hydro power in Norway from the Norwegian Water Resources and Energy Directorate as well as data on electricity generation from wind for Germany obtained from the database of the European Network of Transmission System Operators for Electricity (ENTSO-E). None of the time series showed a significant effect in our regression analysis.

We also calculate a fuel switching price from our coal and gas price data. However, since it is a linear combination of the two series, the effect is already captured with the inclusion of the individual price series. We show this in Appendix B.3 where we also present the results from the inclusion of the additional variables hydro and wind power generation. In the main text, we focus on the set of classical abatement-related price drivers: coal and gas prices as well as economic activity and temperature. The time series are plotted in Figure 2. We observe in Panels (b) and (c) that both the gas and the oil price display a similar upward trend as the allowance price at the end of our sample period. As both variables influence the allowance price, their development could be a potential driver of the trend in allowance prices. However, our further analysis shows that this is not the case.

3.2 Stationarity properties

One of the assumptions of the regression analysis we perform in Section 4 is that the data are stationary. Previous research finds that the allowance price series contains a unit root and therefore, the log return series, calculated as $r_{yt} = \ln(y_t/y_{t-1})$, is used whenever stationarity is needed.⁸

To investigate the stationarity properties of our data, we use a battery of unit root tests. Most of them consider a unit root under the null hypothesis versus stationarity around a linear trend under the alternative. Standard unit root tests, e.g. the Augmented-Dickey-Fuller test (Said and Dickey, 1984), which have been applied in the previous literature, consider a fixed mean or a linear trend under the alternative. It has been shown that this can quite frequently lead to spurious rejections of the unit root null hypothesis. Working with data in first differences if the data are trend stationary can substantially change the results. Therefore, it is crucial to carefully select the alternative hypothesis of unit root tests.

Next to the Augmented-Dickey-Fuller (ADF) test with a linear trend, we apply the tests by Phillips and Perron (1988), Leybourne et al. (1998) and Kwiatkowski et al. (1992). While the Phillips-Perron (PP) test has very similar properties, the test by Leybourne et al. (1998) (LNV) considers a smoothly varying time trend under the alternative hypothesis. The trend can undergo one transition and the time point as well as speed of the transition is determined endogenously by the test. This modification allows for much more flexibility under the alternative which is suitable for our complex data series. The KPSS test by Kwiatkowski et al. (1992) is the only test that considers (trend) stationarity under the null hypothesis and

⁸In the remainder of the paper we often leave out the term "log" and use the terms "log returns" and "returns" interchangeably.

a unit root process under the alternative.

Table 1: Unit root tests

	ADF test		PP test	
	y_t	r_{y_t}	y_t	r_{y_t}
EUA	-0.56	-14.13	-0.68	-17.64
Coal	-2.24	-12.88	-2.09	-16.49
Gas	-2.00	-13.59	-2.11	-18.68
Oil	-1.78	-13.88	-1.76	-18.59
Stoxx 50	-3.42	-15.97	-4.34	-19.11
Stoxx 600	-3.45	-15.90	-4.05	-19.87
Temp	-13.05	–	-15.03	–
Critical values: (-3.13, -3.42, -3.98)				
	LNV test		KPSS test	
	y_t	r_{y_t}	y_t	r_{y_t}
EUA	-1.01	-14.34	1.31	0.07
Coal	-2.34	-13.05	0.55	0.06
Gas	-2.03	13.72	0.57	0.07
Oil	-2.033	-13.96	1.13	0.08
Stoxx 50	-3.71	-16.21	0.74	0.10
Stoxx 600	-3.50	-16.30	0.45	0.07
Temp	-13.15	–	0.07	–
(-4.55, -4.83, -5.42)			(0.12, 0.15, 0.22)	

Source: Own calculations.

Notes: Results from four different unit root tests. The test statistics are given for the data in levels (y_t) and in log returns (r_{y_t}). Below each test are the 90%, 95% and 99% critical values in brackets. The critical values for the ADF and PP tests coincide. The tests were performed with the Bayesian Information Criterion for lag length selection and a maximum number of 8 considered lags. The chosen lag length equals 1 in all cases.

The results are presented in Table 1. The ADF test results are presented in the first two columns. They show that in most cases, the unit root null hypothesis cannot be rejected at a 1% significance level when we look at the data in levels. Exceptions are the temperature series and the two stock

indices (at a 10% level for STOXX 50 and at a 5% level for STOXX 600). The return series are all stationary. The PP test which is given in the next two columns are very similar. The unit root null hypothesis is now also rejected for the two stock indices at the 1% significance level. Results of the more flexible LNV are given in the next two columns. Again, all series but the temperature data contain a unit root. The KPSS test, which is presented in the last two columns, comes to the same conclusion.

Given these results, we conclude that all fuel prices and stock indices contain a unit root and the temperature data are stationary.

3.3 Empirical strategy

Our approach consists of combining three statistical methods, which we apply consecutively to obtain a reliable basis of empirical evidence before drawing conclusions and policy implications. In this section, we briefly outline the three different steps and the rationale for choosing them.

In the first step, we revisit the fundamental price drivers. This is an important step since a tighter link to market fundamentals combined with the increasing gas prices and economic activity (as observed in Section 3.1) offers a potential explanation for the upward trend in EUA prices. The previous literature, which we reviewed in Section 2, does not find a strong link between allowance prices and market fundamentals overall. The gas price, however, has the strongest influence on allowance prices across the reviewed studies. We repeat the linear regression exercise performed by many of the previous studies to determine if gas prices also have a significant effect in our longer data set. In addition, higher (expected) scarcity of allowances due to the reform can lead to a tighter link between allowance prices and market

fundamentals which would only be visible in the final parts of the data. To be able to investigate this properly, we apply a time-varying regression approach which offers a way to flexibly model relationships between multiple time series. A technical difficulty lies in obtaining reliable confidence intervals to judge the significance of the resulting estimates given that our data is serially correlated and show a time-varying variance. We overcome this problem by applying a specific bootstrap method, the autoregressive wild bootstrap, which is robust to serial correlation and heteroskedasticity. The results indicate that the recent price increase is not picked up by any of the price drivers and can therefore not be explained by the increase in economic activity or gas prices. There is also no indication of a tighter relationship between allowance prices and market fundamentals. The regression model as well as the detailed results are presented in Section 4, while the technical details of all methods are given in Appendix A.

To further investigate the upward trend, we apply a statistical test for explosive behaviour in the second step of our analysis. The method of our choice is taken from Phillips et al. (2015) whose approach is based on a right-sided unit root test. Their method has gained wide popularity in the finance literature since its introduction. If the test detects explosive behaviour, Phillips et al. (2015) propose a datestamping procedure to locate the start and end point of the explosive episode. Again, to make the results robust to serial correlation, we apply a bootstrap method as proposed by Pedersen and Montes Schütte (2017). Our particular interest is to find out whether such a period can be detected, and if so, to locate the start and end point. This enables us to compare the onset of the explosive period with the time of adoption of the reform and we can further investigate if it

persists until prices have reached their peak value, as theory would suggest.

Finally, having gained statistical evidence for the existence and location of an explosive period, we model the episode using noncausal processes. Noncausal autoregressive processes were recently proposed by Gouriéroux and Zakoïan (2017) to model speculative bubbles in financial markets because of their ability to mimic locally explosive patterns in time series data. A rapidly emerging econometric literature has been applying these models to diverse bubble phenomena appearing for instance in time series of inflation rates (Hecq and Sun, 2019), stock indexes (Gouriéroux and Zakoïan, 2017; Fries, 2018a), cryptocurrency exchange rates (Hencic and Gouriéroux, 2015; Cavaliere et al., 2020), volatilities (Hecq et al., 2016), and commodity prices (Fries and Zakoïan, 2019; Voisin and Hecq, 2019; Hecq and Voisin, 2019). We leverage the analytical expression of the crash odds of explosive episodes obtained by Fries (2018a) in the framework of noncausal processes to propose a quantification of the (un)sustainability of the recent explosive trend in EU ETS prices.

4 Fundamental price drivers – revisited

If the price run up reflects an increasing scarcity of allowances, fundamentals should have gained importance as price drivers. Accordingly, coefficients would have increased gradually, which we examine using a linear regression model with time-varying coefficients. Such models are flexible, and allow uncovering changes in the relationship without restrictive assumptions regarding the form of the change. We analyse the full time series and not simply the period of the price run up between 2017 and 2018. In fact, we find significant changes in coefficients in earlier years that

have not been identified in the previous literature.

The longer time span of our dataset compared to previous studies provides us with an opportunity for a better understanding of the relationship between allowance prices and abatement-related price drivers. In addition, to be able to compare our results to previous findings, we also run a linear regression model with constant coefficients.

Both approaches, constant and time-varying coefficients, can be captured using the following linear regression framework. For $t = 1, \dots, T$, let

$$r_{EUA,t} = \beta_{0,t} + \beta_{1,t}x_{1,t} + \beta_{2,t}x_{2,t} + \dots + \beta_{m,t}x_{m,t} + \epsilon_t = \boldsymbol{\beta}'_t \mathbf{x}_t + \epsilon_t \quad (4.1)$$

be a model for the return on allowance prices $r_{EUA,t}$, where

$\boldsymbol{\beta}_t = (\beta_{0,t}, \beta_{1,t}, \dots, \beta_{m,t})'$ is a vector of coefficients and $\mathbf{x}_t = (1, x_{1,t}, \dots, x_{m,t})'$ is a set of potential (stationary) price drivers. The error term ϵ_t may be correlated and may exhibit changing variance over time. This simple framework allows for constant coefficients by imposing the restrictions $\beta_{j,t} = \beta_j$, for $j = 0, \dots, m$ and for all $t = 1, \dots, T$. We estimate this linear regression specification by ordinary least squares (OLS).

Without restrictions on the coefficients, Equation (4.1) represents a time-varying coefficient model including a deterministic time trend $\beta_{0,t}$ as well as covariates with time-varying coefficient functions $\beta_{j,t}$ for $j = 1, \dots, m$. All coefficients are assumed to be functions of time which need to follow certain smoothness conditions.⁹

⁹The assumptions require that, for instance, the functions are twice continuously differentiable and bounded. We refer to Cai (2007) for the complete set of assumptions. In addition, as a necessary step to ensure the consistent estimation of these functions, they have to be defined on the interval $(0, 1)$. Therefore, we need to map all points to this interval such that, formally, we have $\beta_{j,t} = \beta_j(t/T)$. This is explained, for instance,

A key advantage of the time-varying coefficient model is that the form of the coefficient functions does not have to be specified in advance. It lets the data determine the form of the relationship. Due to this nonparametric nature the model offers a large degree of flexibility and generality. Estimation of such models is well-established in the econometric literature. We use the local linear kernel estimator of Cai (2007). The estimator fits a locally weighted least squares regression to a neighborhood around each time point in the sample. The weighting function is called the kernel function. It has been shown that the estimator is known to reduce the boundary effects which are a concern in nonparametric estimation. For more information on the estimator and its properties, we refer the interested reader to Fan and Gijbels (1996) for a general overview, and to Cai (2007) for details on the estimator in the context of model (4.1).

To be able to judge the significance of our results, we construct 95% confidence intervals around the nonparametric estimates using a bootstrap method. We rely on the autoregressive wild bootstrap which offers robustness to serial correlation as well as heteroskedasticity. Using nonparametric estimation with bootstrapping is a powerful combination which has been studied in the econometric and statistical literature by e.g. Bühlmann (1998), Neumann and Polzehl (1998) and Friedrich et al. (2019). Details on the estimation technique and the bootstrap algorithm are given in the Technical Appendix. It is particularly suitable for applications in the EU ETS market as in previous studies the residuals of the model experienced heteroskedasticity (e.g. Koch et al., 2014; Lutz et al., 2013). An additional advantage of the method is that it can also be applied when data points

in Robinson (1989).

are missing such that there is no need to resort to interpolation techniques when some data series are incomplete. Although missing data are not a major concern in our application, it nevertheless occasionally happens that single data points are missing from some series. In this case, we simply delete the corresponding data points from our dataset. All empirical results in this section are obtained using the stationary data as determined in Section 3.2.

Both model specifications of (4.1) are not designed to explain sudden jumps in allowance prices. To investigate whether this is a serious issue in our application, we remove outliers and subsequently re-estimate the model.¹⁰ For this we apply the impulse indicator saturation (IIS) approach proposed in Santos et al. (2008). With this method, conditioning on our set of explanatory variables, we find seven outliers in the EUA return series which are removed. More information on the location of the outliers and the IIS application can be found in Appendix B.1.

4.1 Linear regression results

To get a first understanding of the data and to be able to compare our results to the previous literature, we obtain a set of results from Model (4.1) with constant coefficients. Table 2 displays the results from the corresponding OLS regression. We present the coefficient estimates with corresponding Newey-West standard errors (Newey and West, 1987), which are robust to mild forms of autocorrelation and heteroskedasticity. We consider three specifications, where we include the oil price, the STOXX 50 index and the STOXX 600 index, respectively. In specification (a), the

¹⁰Results for the original data (before the removal of outliers) are very similar and they can be found in Appendix B.1.

Table 2: Linear Regression Results

	(a)			(b)		
	$\hat{\beta}_j$	se_{NW}	p -value	$\hat{\beta}_j$	se_{NW}	p -value
Coal	-0.119	0.094	0.206	-0.061	0.097	0.528
Gas	0.190	0.075	0.012	0.198	0.074	0.007
Oil	0.214	0.069	0.002	–	–	–
Temp	-0.001	0.001	0.572	-0.001	0.001	0.450
Stoxx 50	–	–	–	0.139	0.103	0.031
Stoxx 600	–	–	–	–	–	–

	(c)		
	$\hat{\beta}_j$	se_{NW}	p -value
Coal	-0.07	0.097	0.425
Gas	0.198	0.074	0.008
Oil	–	–	–
Temp	-0.001	0.001	0.451
Stoxx 50	–	–	–
Stoxx 600	0.296	0.110	0.007

Source: Own calculations using R.

Notes: Results obtained using OLS estimation. The dependent variable is the return on EUAs and the set of (stationary) regressors changes in each specification. The p -values are based on Newey-West standard errors.

two significant factors are the gas and the oil price. The coal price, as an important driver, does not show a significant effect on the allowance price in this initial regression. The estimates are very similar in specifications (b) and (c) where we include the stock indices.

This does not come as a surprise given the results from previous studies, as we do not split the data into sub-periods, nor do we include any dummy variables to take out the effect of major policy announcements. In all specifications, both significant coefficients are positive and thus show the anticipated sign. The coefficient of the coal price also shows the sign

predicted by economic theory, while being insignificant.¹¹

4.2 Time-varying coefficient results

We now apply the nonparametric approach to our data. The results presented in this section are obtained using the gas, coal and oil price returns as explanatory variables as well as the temperature series after seasonality has been removed. As a robustness analysis, we provide additional results (e.g. from the inclusion of the other variables) in Appendix B.3.

The estimated trend and coefficient curves (blue) together with their 95% confidence intervals (orange) can be found in Figure 3. As with linear regression, a coefficient is significant if zero (indicated by the gray line) does not fall within the confidence interval. The main difference to parametric regression is that there is not merely one coefficient whose estimate and confidence interval permits a verdict on the question of significance of an explanatory variable. We have a coefficient estimate and a corresponding confidence interval for each point of the sample. Thus, there can be periods of significance and insignificance as well as changes in the sign and magnitude of a coefficient.

Against this background, we observe in Figure 3 that the significance of all of the included variables changes over time. For all variables, there are period with a significant effect as well as periods of insignificance. All graphs further have in common that the width of the confidence intervals changes over time and in many cases, become substantially wider at the beginning and the end of the sample. The widening at the beginning and the end is common in the application of the bootstrap method and reflects

¹¹In Appendix B.3 we include each stock index in addition to the oil price and show that they do not have a significant effect.

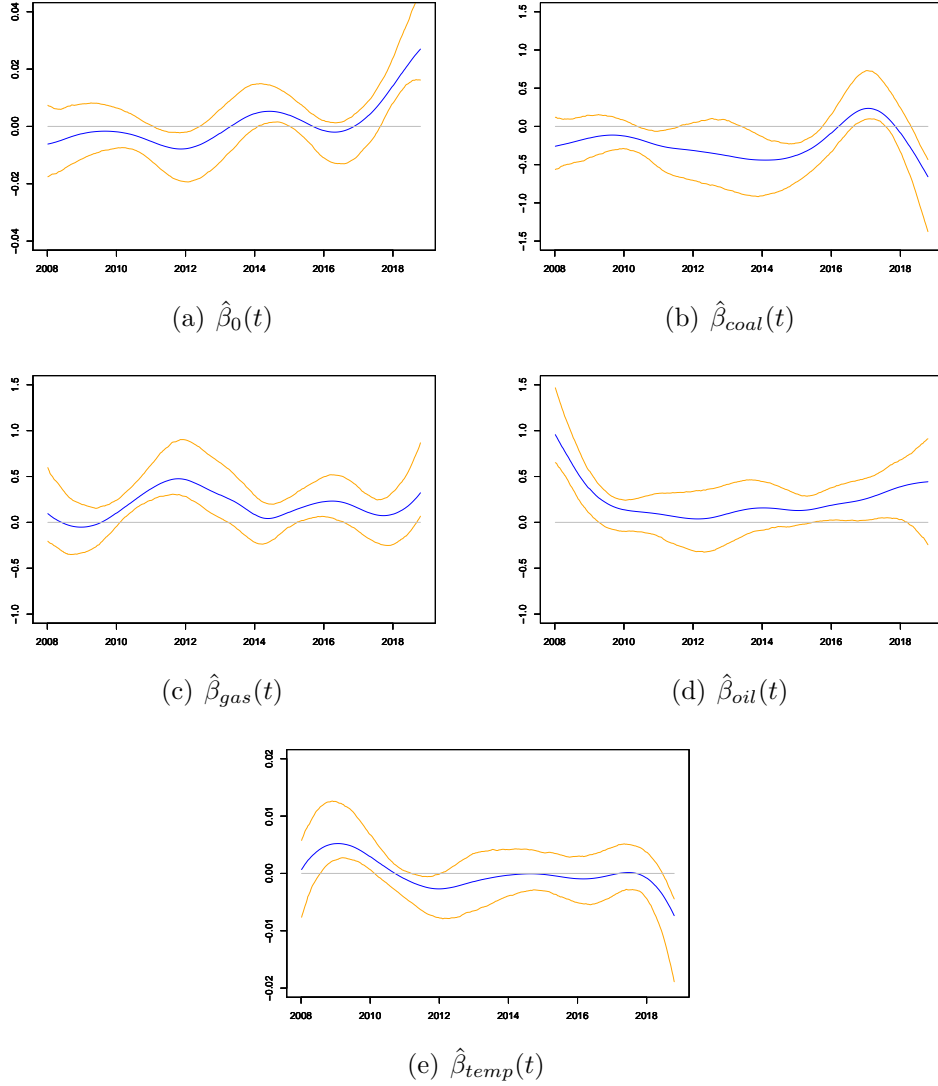


Figure 3: Time-varying coefficient estimates

Source: Own calculations.

Notes: Nonparametrically estimated coefficient curves (blue) and 95% confidence intervals (orange) from a regression of the return on EUA on a trend and the four main price drivers. For the nonparametric estimation, we apply the local linear estimator with the Epanechnikov kernel given by the function $K(x) = \frac{3}{4}(1-x^2)\mathbb{1}_{\{|x|\leq 1\}}$. We use a bandwidth parameter of $h = 0.09$ which we justify in Appendix B.2. The bootstrap procedure is applied using 999 replications.

that the first and last estimates can be inaccurate. As the nonparametric estimator uses a two-sided window for estimation, there are boundary effects which can distort the first and last 20 estimated points on the coef-

ficient curves. This contributes to the widening of the confidence intervals at these points. Widening in the middle of the sample can be an indication of heteroskedasticity or autocorrelation in the data over this time span. In such cases, the bootstrap method accounts for these irregularities by making the confidence intervals wider such that the nominal confidence level (in this case of 95%) can be maintained.

In Panel (a), the nonparametric trend fluctuates around zero for most of the considered time span. However, at the end of the sample, it turns significantly positive. This shows that the recent price increase in allowance prices seems to be picked up by the trend component. Panel (b) shows that the coal price has a significant negative effect for two periods: one ranging from 2010 to mid 2011 and one from 2013 to mid 2015. Subsequently, the coefficient becomes positive and significant for a short period in 2016/2017. This is a very interesting finding given that its coefficient was found to be insignificant in the linear regression analysis. This suggests that the time variation might have caused the insignificance in the linear regression results. Moving on to Panel (c), we see the coefficient of the gas price series. It has the expected positive sign and is significant over long periods of the sample. This is in line with previous findings as well as the linear regression results presented in Section 4.1. However, we also find a period of insignificance in the usually stable gas price coefficient which is located in 2014. The coefficient of the oil price, as displayed in Panel (d) is positive and significant over two periods – until 2009 and from 2015 onward. Finally, from Panel (e) we see that the temperature series shows only two very short periods of significance and the magnitude of the coefficient estimate is, as in the linear regression results, negligible.

Overall, our analysis provides evidence of time variation in the relationship between allowance prices and the considered price drivers. It stresses the need that the relationship has to be modeled with care. Due to the various periods of insignificance, our results offer a potential explanation for insignificant coefficients found with linear regression techniques used in some of the previous work. In addition, although the method offers great flexibility in modelling the relationship, the recent upward trend stays in the trend component. This implies that fundamentals have not gained relevance as drivers during the price run up. We now further examine the recent price increase using a combination of statistical tests and models potentially explosive behaviour.

5 Explosive behaviour in allowance prices

In the previous section we saw that the recent upward movement was not picked up by any of the included price drivers and one might suspect that the recent trend could represent a period of explosiveness. It is therefore our goal in this section to further investigate this recent price development. Our method involves two steps. First, we apply the approach by Phillips et al. (2015) which provides a refined way to test for periods of explosive autoregressive behaviour and, in an additional step, to locate them.¹² The test is based on a right-sided unit root test performed on a forward and backward expanding window. We indeed find a significant explosive period in the allowance price series which overlaps with the recent price increase. In a second step, we further analyse the explosive period we found in the

¹²While *explosive autoregressive behaviour* is the most accurate way of describing the phenomenon we test for, we use this terminology interchangeably with *explosive behaviour*, *explosiveness* and sometimes *bubble* in the remainder of this paper. We elaborate further why we have to be careful with the term *bubble* in Sections 5.2 and 6.

first step. With the help of a novel class of statistical models we predict ex-ante probabilities of collapse of the explosive period in allowance prices which we compare to the out-of-sample price behaviour.

5.1 Testing for explosive episodes

To investigate whether the upward trend in allowance prices constitutes an episode of explosive behaviour, we apply the testing procedure of Phillips et al. (2015). Similarly to ordinary unit root tests, it tests the null hypothesis of a unit root in the data. The construction of the test is based on the ADF test as performed in Section 3.2, but it considers a different alternative hypothesis. Instead of stationarity, the Phillips et al. (2015) test considers mildly explosive behaviour under the alternative. An additional difference is that the test is applied to sub-samples of the data. For rejection of a unit root in favor of explosive episodes it is sufficient if the data show such behaviour in a sub-sample. As the name explosive episode or bubble suggests, this is not a long-term phenomenon. The test applies an ADF test to various subsamples. Subsequently, the test statistic is constructed as the supremum over all ADF test statistics. It is therefore called Generalised Supremum ADF (GSADF) test. In case of rejection, there is evidence of explosive behaviour. For such cases, Phillips et al. (2015) additionally propose a method to locate the beginning and end point of the episode. Details on the testing procedure and the date stamping are available in our Technical Appendix.

A bubble is present where prices diverge from fundamentals. In this context, it means that there is evidence for an explosive period when the price series shows such behaviour but the fundamental drivers do not. Therefore,

Table 3: GSADF test results

Test statistics		Critical values (90%, 95%, 99%)	
Variable	GSADF	simulated	bootstrap
EUA	3.998		(2.270, 2.555, 3.201)
Coal	1.676		(2.487, 2.795, 3.446)
Gas	1.299	(1.983, 2.175, 2.608)	(2.383, 2.645, 3.372)
Oil	2.722		(2.200, 2.505, 3.104)
Stoxx 50	0.782		(2.310, 2.668, 3.099)
Stoxx 600	0.953		(2.302, 2.302, 3.170)

Source: Own calculations using Matlab and the R package MultipleBubbles.

Notes: The GSADF test statistics with simulated critical values (2000 repetitions) and bootstrapped critical values (4999 repetitions). The test is applied using a minimum window size of 40 observations, which is chosen according to the rule $r_0 = 0.01 + 1.8\sqrt{T}$ suggested by Phillips et al. (2015). For the selection of the number of lags we use the Bayesian Information Criterion.

we also apply the tests to the main price drivers. We test coal, gas and oil prices as well as the two stock indices.¹³ The test was proposed under the assumption of uncorrelated and homoskedastic error terms. It has been shown in e.g. Harvey et al. (2016) and Pedersen and Montes Schütte (2017) that a violation of this assumption leads to overrejection which can be resolved by a bootstrapping approach. In our application, we rely on both, the simulated critical values and the bootstrap approach of Pedersen and Montes Schütte (2017).¹⁴

Table 3 presents the results; the left part of the table shows the test statistics for each considered time series and the right part gives the two


¹³We do not consider temperature data here given the fact that it is a climatological time series which due to its natural properties will only experience gradual structural change rather than explosiveness (this is also confirmed by the plot of the data in Figure 2(f)).

¹⁴We thank the authors for kindly providing us with code for the bootstrap test.


different sets of critical values – the critical values computed by Monte Carlo simulation as described in Phillips et al. (2015) and the bootstrapped critical values. For both tests, we would like to use significance levels of 10%, 5% and 1%. As the test is right-sided, this can be achieved by the 90%, 95% and 99% critical values, respectively. Note that, as the bootstrap results will be conditional on the original sample, we have a slightly different set of critical values for each series.

Both tests come to the same conclusion. They present evidence of explosive behaviour not only in the allowance price series but also in the oil price series. For the coal and gas prices and the stock indices, the null hypothesis of a unit root is not rejected and hence, these factors are excluded as possible drivers of the movement. In contrast, the oil price cannot be excluded. Therefore, for the two cases of rejection we move to the date stamping procedure, which is determined by the BSADF test sequence which is short for Backward Supremum ADF test. We again obtain critical values from simulation and bootstrapping using the same specifications as for the GSADF test. The results using bootstrapped critical values are presented in Figure 4. It simultaneously plots the series of critical values (orange) and the test statistics (blue). Panel (a) gives results for the allowance price series and Panel (b) for the oil prices. Using simulated critical values leads to very similar conclusions. The results can be found in Appendix B.

There is clear evidence for a long period of explosive behaviour in allowance prices, starting at the beginning of 2018. The series of test statistics starts exceeding the critical value series in early March of 2018 and does not cross it again until the end of the sample. Before 2018, we can see



(a) EUA prices



(b) Oil prices

Figure 4: Date stamping of explosive episodes

Source: Own calculations.

Notes: Results from the BSADF test as proposed by Phillips et al. (2015), applied to the EUA prices in Panel (a) and the oil prices in Panel (b). Each panel shows the bootstrapped critical value series (orange) and the test statistics (blue).

several such episodes in Panel (a) with the most pronounced one located in 2016. Starting in mid January 2016, the BSADF statistic lies above the critical value sequence for 11 weeks until the beginning of April 2016. This exceeds the established minimum duration of 7 weeks. Observing the allowance price development at the time, we see that this period is caused by a price drop in spring 2016.¹⁵ The spikes in the test statistic before 2016

¹⁵Some analysts have portrayed the price drop as a potential consequence of the COP15 outcome. But to the best of our knowledge none of the scientific event studies on EUA prices has covered this period.

do not last longer than 3 weeks and therefore, cannot be taken as evidence for explosive behaviour.

In 2018, the length of the explosive episode clearly exceeds the minimum duration. Comparing this to the oil price results in Panel (b), we do not find overlapping explosive behaviour in this series, although the test detects a potential explosive period and we observe an upward trend in oil prices at the end of the sample. The period in oil prices which caused the GSADF test to reject is located in 2014. Starting mid 2014 oil prices fell for a period of more than 6 months indicating that it was a period of collapse and not exuberance that is picked up by the test.

Specifically, the period of collapse is first detected by the test in mid October 2014 where the test statistic lies above the critical values for three weeks, then drops for one week before it exceeds it again for 11 consecutive weeks. This period lasts from November 2014 until the beginning of February 2015. This is in line with the findings in Sharma and Escobari (2018) who investigate several different oil price series. The authors find a significant period of collapse with similar timing in all investigated oil prices. In addition, the absence of evidence for periods of exuberance or collapse in the gas price series is also confirmed by Sharma and Escobari (2018).¹⁶

Combining these results provides statistical evidence that the allowance price experienced an explosive period which is not accompanied by similar behaviour in any other series we considered. Although the gas and the oil price experienced a simultaneous upward trend at the end of our sample period, the formal test allows us to exclude both factors as potential drivers

¹⁶For completeness, we also apply the GSADF test to the fuel switching price. Unsurprisingly, we do not find evidence for explosive behaviour in this series.

of the rapid EUA price increase. This is to our knowledge a first empirical result pointing in the direction of a period of exuberance in EUA prices which is not driven by abatement-related fundamentals. While this finding strongly suggests the existence of a price bubble, it should be noted that there could be other explanations which are in line with these results. For instance, non-explosive changes in fundamentals could lead to new equilibrium price levels and the path to the new levels can potentially appear explosive. This has been pointed out by Harvey et al. (2016). As shown in Section 4, changes pointing towards an increasing role of fundamentals as a result of the reform could not yet be found in the data.

5.2 Predicting probabilities of collapse

The statistical evidence points to an explosive episode whose drivers are not related to fundamentals. Non-fundamental bubbles in prices are often seen as a result of widely shared exuberant beliefs among market participants as explained in the Introduction. Recently, rational expectation specifications of financial bubbles have been shown to exhibit as particular solutions so-called noncausal, or non-fundamental, autoregressive time series models (Gouriéroux et al., 2020), that is, autoregressive processes with roots located inside the unit circle. These statistical models, compatible with rational expectation price models, have the ability to parsimoniously mimic explosive episodes in financial time series data and are increasingly applied in the economic and statistical literature to analyse time series with forward-looking components.

Fries (2018a,b) derives analytical expressions for the ex ante crash odds of bubbles generated by noncausal models driven by extreme, power-law

distributed shocks, which we exploit here to quantify the (un)sustainability of the recent explosive episode in the EUA price. We focus on the arguably simplest noncausal model – the noncausal AR(1) – which mimics bubbles with an exponentially-shaped inflation phase culminating in a peak before collapsing to pre-bubble levels. There are two main reasons for this choice: (i) the obtained exponential fit of the recent upward trend in the EUA prices appears excellent, (ii) bubbles generated by the noncausal AR(1) feature a memory-less, or non-aging, property which implies that the crash date cannot be known with certainty by market participants and indicates compatibility of the model with a no-arbitrage setting. Further details on the approach are provided in the Technical Appendix.

Modelling the explosive episode in the EUA price as an ongoing realisation of a noncausal AR(1) bubble, the sustainability of the upward trend can be completely characterised by the given of two parameters: the growth rate of the exponential inflation phase, and a parameter describing the likely height of the incoming peak – the tail exponent of the shocks' power-law distribution. A higher growth rate and a likely lower peak (i.e. higher tail exponent) entail a less sustainable explosive episode. The growth rate can be conveniently obtained by fitting an exponential trend on the recent data – which boils down to a linear regression of the log-prices on time. To this end, we include all data points post-March 2018 into the exercise, as this date marks our earliest statistical evidence of explosive episode based on Phillips et al. (2015) test. For sensitivity analysis, we also consider the estimate of the growth rate using the data points since the beginning of the upward trend, mid 2017. The second parameter, the power-law tail exponent, is more elusive and cannot be easily retrieved

from the data. We prefer to remain as agnostic as possible regarding the value of this second parameter, and only presume it lies within the widest range of reasonable and admissible values under Fries (2018a) framework.¹⁷

Table 4 gathers the results and Figure 5 presents the exponential fits

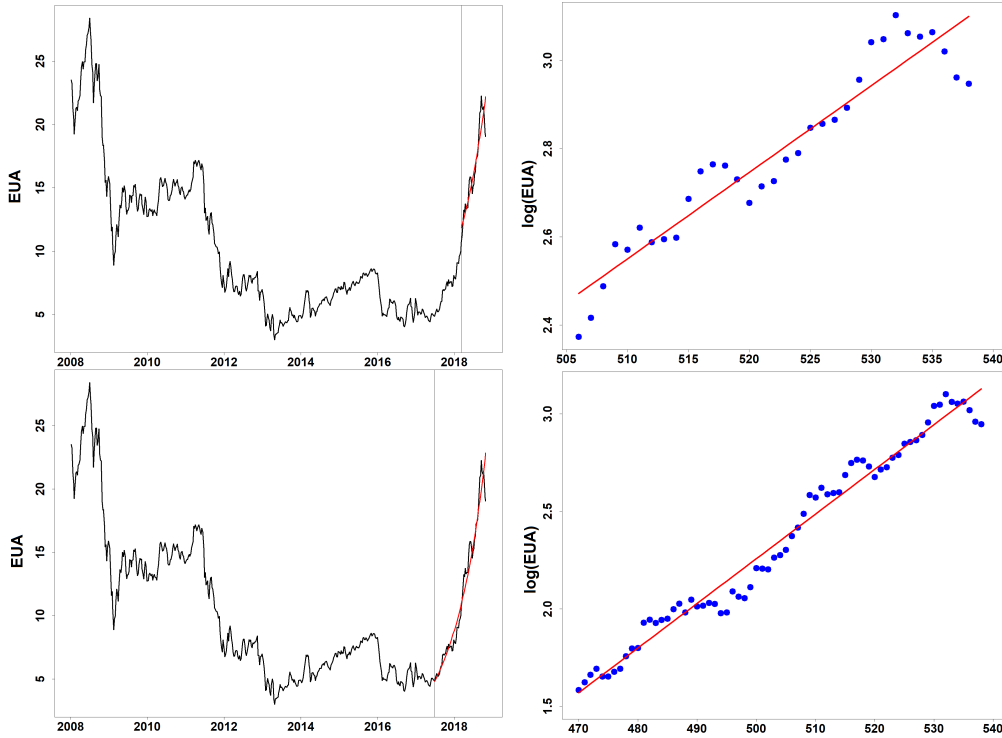


Figure 5: Exponential trend fitted on the EUA prices (left panels) based on the linear regressions of log-EUA prices on time and intercept $\log(\text{EUA}_t) = at + b$ (right panels). Upper panels: fit obtained using the data points post-March 2018 ; Lower panels: fit obtained using the data points post-June 2017.

of the recent EUA data. The obtained growth rates using either the data post-March 2018 or post-June 2017 both point to a similar pace of around

¹⁷ This range encompasses values of the power-law tail exponent, denoted α , from $\alpha = 0.5$ to α arbitrarily close to 2: that is, from very heavy-tailed shocks as extreme as a Lévy distribution which have even infinite expectation, to much milder shocks close to have finite variance. Most studies in the financial literature report values from significantly below one up to four for financial series (Ibragimov and Prokhorov (2016) and the references therein), with reported values above two furthermore not necessarily evidence against the infinite variance hypothesis (McCulloch, 1997), which is used in Fries (2018a,b).

	Based on EUA data from 03/2018 to 10/2018		Based on EUA data from 06/2017 to 10/2018	
Exponential trend monthly growth rate \hat{a}	8.1%		9.5%	
Noncausal AR coefficient $\hat{\rho}$	0.92		0.91	
Plausible range for tail exponent α	[0.5 – 2]		[0.5 – 2]	
Odds of crash within m months	Lower est.	Upper est.	Lower est.	Upper est.
$m = 1$	4.0%	15%	4.6%	17%
3	11%	39%	13%	44%
6	22%	62%	25%	68%
12	39%	86%	44%	90%
Expected explosive episode duration (in months)	7	25	6	22

Table 4: Estimated monthly growth rates \hat{a} of exponential trends fitted on the EUA log-prices $\ln(\text{EUA}_t) = at + b$, based on data since March 2018 (left) and since mid-2017 (right) ; Corresponding noncausal AR(1) coefficients $\hat{\rho} = \exp(-\hat{a})$; Uniform prior on power-law tail exponent $\alpha \in [0.5, 2]$; Corresponding lower and upper estimates for crash odds within $m = 1, 3, 6, 12$ months computed as $1 - (1/2)^{m/\hat{h}_{0.5}}$, where $\hat{h}_{0.5} = -\frac{\ln 2}{\alpha \ln(\hat{\rho})}$ is the half-life of an α -stable noncausal AR(1) bubble with parameters $\hat{\rho}$ and α ; Expected duration of the explosive episode computed as $\frac{1}{1-\hat{\rho}^\alpha}$.

8-9% price growth per month. Moreover, the fits on the log-prices appear at least adequate on the data post-March 2018 and even excellent on the data post-June 2017. This comforts the assumption of an exponentially-shaped bubble *à la* noncausal AR(1), which appears very much compatible with the price trajectory. Because we assume a range of possible values for the tail exponent parameter rather than a specific value (see Footnote 17), we obtain lower and upper estimates of the odds of a crash occurring at future horizons. We report these lower and upper estimates for 1, 3, 6 and 12 months horizons beyond the end of our sample. The crash odds that we

find are relatively important: already at 1 month horizon, the likelihood of a crash ranges from 4% (lower estimate) up to 17% (upper estimate). According to our lower estimates, the likelihood of a crash occurring within 12 months are high but nevertheless below 50%, whereas our upper estimates appear to place the 50% chance tipping point much earlier, between 3 and 6 months horizon. Another equivalent measure of sustainability of an explosive episode (in the framework of the assumed model) is that of its expected duration. According to our same estimates, the expected duration of the EUA explosive episode should range from 6 (lower est.) to 25 months (upper est.).

A posteriori, in 2020, it is apparent on updated data that the explosive trend wore out very quickly after the end of our sample. Even though an abrupt vertical crash down to pre-bubble level did not occur, the ex-post explosive trend (un)sustainability appears highly compatible with our estimates. The fact that the exponential trend leveled off and yet the prices did not collapse could be an indication that the drivers of explosive growth vanished. It is, however, unclear whether this implies a long-term stabilization of prices or a postponement of the collapse.

6 Discussion and conclusion

Results of the first two methods clearly suggest that the reform has triggered market participants into speculation about the reform's price effects. To begin with, results from the time-varying coefficient model confirm that the recent upward trend cannot be explained by movements in the considered price drivers. Instead it is picked up by the trend component, which implies other price drivers. Furthermore, we find evidence for market over-

reaction indicated by an explosive period beginning in March 2018 – its onset thus exactly coincides with the month the reform has been passed into law. This event did not come as a complete surprise to market participants: already in November 2017, the European Council endorsed the reform proposal and submitted it to the European Parliament for approval.¹⁸ However, our results suggest that the market reacted more strongly to actual legislation rather than to mere news about it, particularly in terms of new players deciding whether to enter a market or not (see Section 1). This suggests that sentiment about the future scarcity of allowances may have indeed been the decisive driver.

Yet, reform-induced speculation as such does not necessarily imply a destabilisation of the market. It could potentially reflect the search for the new equilibrium price induced by the reform, i.e. an adaption process which, as previously mentioned, can appear explosive (Harvey et al., 2016). Speculation only destabilises the market if it leads to an overreaction – an inflated price that exceeds the new long-term equilibrium level. If that was the case, prices would be bound to collapse in the future. To shed light on this issue, we employ additional statistical techniques to estimate the probabilities of collapse for different time horizons, starting in October 2018 when the price increase peaked. The upper estimates for the one year ahead crash odds are as high as 86%. Looking at how prices evolved after October 2018, no collapse to rock bottom levels materialized. However, a (partial) collapse from around 25 EUR/t down to around 15 EUR/t in November 2018 can be observed. Subsequently, the price rebounded and it

¹⁸<https://www.consilium.europa.eu/en/press/press-releases/2017/11/22/reform-of-the-eu-emissions-trading-system-council-endorses-deal-with-european-parliament/>

is difficult to judge if this drop was a crash or a sign of increasing volatility around a higher (post-reform) equilibrium level.

Deciding between the two scenarios of fluctuations around a higher equilibrium price level or a partial collapse becomes even more difficult once we acknowledge that since the reform, a number of additional announcements have been made towards higher stringency. Already in November 2018, the European Commission published a long-term vision for a climate-neutral economy in 2050.¹⁹ Moreover, in the course of 2019, it became clear that the European Commission would propose more stringent climate targets for 2030, implying a more stringent cap. In December 2019, it eventually published the EU Green Deal which proposes a ratcheting up of long term climate targets.²⁰ This development may have arguably counteracted further price drops. Similarly, subsequent changes in market fundamentals likely had a confounding impact on prices. Analyzing this price impact in detail is beyond the scope of the paper and left for future research.

Although our empirical results lend themselves to multiple interpretations, the characteristics of the reform suggest that it has indeed persistently destabilised the market. As pointed out in the introduction, in an act of "smokescreen politics" regulators leveraged on regulatory complexity to push higher ambition (cancellations of allowances) through the backdoor. The new mechanism to cancel allowances from the MSR is extremely complex, and the volume of allowances that will eventually be canceled from 2024 on is highly difficult to predict. This implies persistent cancellation-related price uncertainty in the coming years. This is enhanced by a feed-

¹⁹<https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52018DC0773&from=EN>

²⁰https://ec.europa.eu/info/publications/communication-european-green-deal_en

back effect on prices: anticipation of cancellation affects current prices, which in turn affects cancellation (Osorio et al., 2020; Quemin, 2020). The broad range of numerical estimates for the number of allowances that will be cancelled compiled by Osorio et al. (2020) clearly underlines this.

Finally, we would like to discuss the related policy implications of our paper. What does this imply for the question of rules vs. discretion? The 2018 reform of the EU ETS defies a clear-cut classification. The actual mechanisms implemented through the reform, such as MSR intake, outtake and cancellation, are rule based. They are generally designed to preclude further discretionary intervention in the future. At the same time, the reform itself is discretionary in the sense that policy makers intervened in the market retrospectively in order to, in their own words, reduce the surplus of emission allowances in the carbon market. Moreover, the complexity of the rules – or rather the difficulty of predicting their impact on prices – tend to increase price uncertainty in the market. This makes it likely that further interventions will eventually have to follow. In fact, as part of the original 2018 reform, a review of the MSR and potential subsequent reform is scheduled for 2021. Hence, regulatory uncertainty inevitably prevails, conflating the post-reform equilibrium price with the uncertainty of how the upcoming MSR revision will affect the market.

How can EU policy makers do better in the future? To answer this question it needs to be acknowledged that, understandably, governments want to have the flexibility to cope with future circumstances (Hasegawa and Salant, 2014). In face of the trade-off between flexibility and commitment, it is crucial to choose the right time horizon that best balances benefits and costs (Hepburn, 2006). Other work stresses the role of process, i.e. how

policy is executed, revised and reformed: Aldy (2020) proposes the institutionalization of an act-learn-act approach for a potential US carbon tax with a view on promoting its adaptability to changing environmental and economic conditions. Correspondingly, for the Clean Air act, Carlson and Burtraw (2019) conclude that if the policy change process is well formalized, predictable and transparent, this can enable policy to adapt to new information while ensuring its durability. While the ETS reform arguably meets the first and second criteria, the "smokescreen politics" obviously failed on the third by implementing mechanisms whose impact on prices is complex and hard to predict.

Consequently, the key to doing better in the future is to simplify the rules in terms of minimizing the price uncertainty they induce. The straightforward option for this seems to be to adapt the cap. However, concerns over intertemporal price formation may need to be addressed in addition. First, policy makers want to contain price volatility, for instance because they are concerned about the political acceptability of high prices (Borenstein et al., 2019). Second, changes in the cap are necessarily long-term and rely on commitment. If this commitment is not credible to market participants, then prices will be distorted (Quemin and Trotignon, 2019). Given these concerns, a price collar could be a helpful complement. It contains price volatility (Burtraw et al., 2010) and effectively works as a commitment device (Fuss et al., 2018). It is simple and transparent since it merely requires defining an upper and lower price path. In summary, cap adjustments as the main lever to control long-term supply in combination with a price collar to enhance commitment and to contain volatility seems to be the right choice – in particular, if suitable adjustment rules can be

found to make the program sufficiently flexible.

References

- Aatola, P., Ollikainen, M., and Toppinen, A. (2013). Price determination in the EU ETS market: Theory and econometric analysis with market fundamentals. *Energy Economics*, 36(October 2007):380–395.
- Abreu, D. and Brunnermeier, M. K. (2003). Bubbles and crashes. *Econometrica*, 71(1):173–204.
- Aldy, J. E. (2020). Carbon Tax Review and Updating: Institutionalizing an Act-Learn-Act Approach to U.S. Climate Policy. *Review of Environmental Economics and Policy*, 14(1):76–94.
- Barberis, N., Greenwood, R., Jin, L., and Shleifer, A. (2018). Extrapolation and bubbles. *Journal of Financial Economics*, 129(2):203–227.
- Borenstein, S., Bushnell, J., Wolak, F. A., and Zaragoza-Watkins, M. (2019). Expecting the Unexpected: Emissions Uncertainty and Environmental Market Design. *American Economic Review*, 109(11):3953–3977.
- Bühlmann, P. (1998). Sieve bootstrap for smoothing in nonstationary time series. *Annals of Statistics*, 26:48–83.
- Burtraw, D., Palmer, K., and Kahn, D. (2010). A symmetric safety valve. *Energy Policy*, 38(9):4921–4932.
- Cai, Z. (2007). Trending time-varying coefficient time series models with serially correlated errors. *Journal of Econometrics* 136, pages 163–188.
- Cambanis, S. and Fakhre-Zakeri, I. (1995). On prediction of heavy-tailed autoregressive sequences: forward versus reversed time. *Theory of Probability & Its Applications*, 39(2):217–233.
- Carlson, A. and Burtraw, D., editors (2019). *Lessons from the Clean Air Act*. Cambridge University Press.

- Cavaliere, G., Nielsen, H. B., and Rahbek, A. (2020). Bootstrapping non-causal autoregressions: with applications to explosive bubble modeling. *Journal of Business & Economic Statistics*, 38(1):55–67.
- Chu, C.-K. and Marron, J. S. (1991). Comparison of two bandwidths selectors with dependent errors. *The Annals of Statistics* 19(4), pages 1906–1918.
- Corbet, S., Lucey, B., and Yarovaya, L. (2018). Datestamping the Bitcoin and Ethereum Bubbles. *Finance Research Letters*, 26:81–88.
- Craven, P. and Wahba, G. (1978). Smoothing noisy data with spline functions. *Numerische Mathematik*, 31(4):377–403.
- Creti, A. and Joëts, M. (2017). Multiple bubbles in the European Union Emission Trading Scheme. *Energy Policy*, 107:119–130.
- Deeney, P., Cummins, M., Dowling, M., and Smeaton, A. F. (2016). Influences from the European parliament on EU emissions prices. *Energy Policy*, 88:561–572.
- Delarue, E., Voorspools, K., and D’haeseleer, W. (2008). Fuel Switching in the Electricity Sector under the EU ETS: Review and Prospective. *Journal of Energy Engineering*, 134(2):40–46.
- Diba, B. T. and Grossman, H. I. (1988). Explosive rational bubbles in stock prices? *The American Economic Review*, 78(3):520–530.
- Fan, J. and Gijbels, I. (1996). *Local polynomial modeling and its applications*. Chapman and Hall, London.
- Friedrich, M., Mauer, E.-M., Pahle, M., and Tietjen, O. (2020). From fundamentals to financial assets: the evolution of understanding price formation in the EU ETS. Technical report, Kiel, Hamburg.
- Friedrich, M., Smeekes, S., and Urbain, J.-P. (2019). Autoregressive wild bootstrap inference for nonparametric trends. *Journal of Econometrics*, In Press.

- Fries, S. (2018a). Conditional moments of noncausal alpha-stable processes and the prediction of bubble crash odds. *ENSAE-CREST and Vrije Universiteit Amsterdam. arXiv preprint arXiv:1805.05397*.
- Fries, S. (2018b). Path prediction of aggregated alpha-stable moving averages using semi-norm representations. *ENSAE-CREST and Vrije Universiteit Amsterdam. arXiv preprint arXiv:1809.03631*.
- Fries, S. and Zakoïan, J.-M. (2019). Mixed causal-noncausal ar processes and the modelling of explosive bubbles. *Econometric Theory*, 35(6):1234–1270.
- Fuss, S., Flachsland, C., Koch, N., Kornek, U., Knopf, B., and Edenhofer, O. (2018). A Framework for Assessing the Performance of Cap-and-Trade Systems: Insights from the European Union Emissions Trading System. *Review of Environmental Economics and Policy*, 12(2):220–241.
- Gouriéroux, C., Jasiak, J., and Monfort, A. (2020). Stationary bubble equilibria in rational expectation models. *Forthcoming in Journal of Econometrics*.
- Gouriéroux, C. and Zakoïan, J.-M. (2017). Local explosion modelling by non-causal process. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(3):737–756.
- Harvey, D. I., Leybourne, S. J., Sollis, R., and Taylor, A. M. R. (2016). Tests for explosive financial bubbles in the presence of non-stationary volatility. *Journal of Empirical Finance*, 38:548–574.
- Hasegawa, M. and Salant, S. (2014). Cap-and-trade programs under delayed compliance: Consequences of interim injections of permits. *Journal of Public Economics*, 119:24–34.
- Hecq, A., Lieb, L., and Telg, S. (2016). Identification of mixed causal-noncausal models in finite samples. *Annals of Economics and Statistics/Annales d'Économie et de Statistique*, (123/124):307–331.
- Hecq, A. and Sun, L. (2019). Identification of noncausal models by quantile autoregressions. *arXiv preprint arXiv:1904.05952*.

- Hecq, A. and Voisin, E. (2019). Predicting bubble bursts in oil prices using mixed causal-noncausal models. *Maastricht University. arXiv preprint arXiv:1911.10916*.
- Hencic, A. and Gouriéroux, C. (2015). Noncausal autoregressive model in application to bitcoin/usd exchange rates. In *Econometrics of risk*, pages 17–40. Springer.
- Hepburn, C. (2006). Regulation by prices, quantities, or both: A review of instrument choice. *Oxford Review of Economic Policy*, 22(2):226–247.
- Hintermann, B. (2010). Allowance price drivers in the first phase of the EU ETS. *Journal of Environmental Economics and Management*, 59(1):43–56.
- Hintermann, B., Peterson, S., and Rickels, W. (2016). Price and market behavior in phase II of the EU ETS: A review of the literature. *Review of Environmental Economics and Policy*, 10(1):108–128.
- Ibragimov, R. and Prokhorov, A. (2016). Heavy tails and copulas: Limits of diversification revisited. *Economics Letters*, 149:102–107.
- Jiménez-Rodríguez, R. (2019). What happens to the relationship between EU allowances prices and stock market indices in Europe? *Energy Economics*, 81:13–24.
- Juhrich, K. (2016). CO2 emission factors for fossil fuels. *Climate Change, Umweltbundesamt*, 28:1–48.
- Klein Tank, A. M. G., Wijngaard, J. B., Können, G. P., 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. F. V., Forland, E., Miletus, 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. V.,

- and Petrovic, P. (2002). Daily dataset of 20th-century surface air temperature and precipitation series for the European Climate Assessment. *International Journal of Climatology*, 22(12):1441–1453.
- Koch, N., Fuss, S., Grosjean, G., and Edenhofer, O. (2014). Causes of the EU ETS price drop: Recession, CDM, renewable policies or a bit of everything?-New evidence. *Energy Policy*, 73:676–685.
- Koch, N., Grosjean, G., Fuss, S., and Edenhofer, O. (2016). Politics matters: Regulatory events as catalysts for price formation under cap-and-trade. *Journal of Environmental Economics and Management*, 78:121–139.
- Kwiatkowski, D., Phillips, P. C. B., Schmidt, P., and Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics*, 54(1):159–178.
- Kydland, F. E. and Prescott, E. C. (1977). Rules Rather than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy*, 85(3):473–491.
- Leybourne, S., Newbold, P., and Vougas, D. (1998). Unit roots and smooth transitions. *Journal of Time Series Analysis*, 19(1):83–97.
- Lutz, B. J., Pigorsch, U., and Rotfuß, W. (2013). Nonlinearity in cap-and-trade systems: The EUA price and its fundamentals. *Energy Economics*, 40:222–232.
- Marcu, A., Alberola, E., Caneill, J., Mazzoni, M., Schleicher, S., Vailles, C., Stoefs, W., Vangenechten, D., and Cecchetti, F. (2019). 2019 State of the EU ETS Report. Technical report.
- McCulloch, J. H. (1997). Measuring tail thickness to estimate the stable index α : A critique. *Journal of Business & Economic Statistics*, 15(1):74–81.
- Neumann, M. H. and Polzehl, J. (1998). Simultaneous bootstrap confidence bands in nonparametric regression. *Journal of Nonparametric Statistics* 9, pages 307–333.

- Newey, W. K. and West, K. D. (1987). A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*, 55(3):703–708.
- Osorio, S., Tietjen, O., Pahle, M., Pietzcker, R. C., and Edenhofer, O. (2020). Reviewing the Market Stability Reserve in light of more ambitious EU ETS emission targets.
- Pedersen, T. and Montes Schütte, E. (2017). Testing for Explosive Bubbles in the Presence of Autocorrelated Innovations. *CREATES Research Paper 2017-9*.
- Perino, G. (2018). New EU ETS Phase 4 rules temporarily puncture waterbed. *Nature Climate Change*, 8(4):262–264.
- Phillips, P. C. B. and Perron, P. (1988). Testing for a Unit Root in Time Series Regression. *Biometrika*, 75(2):335–346.
- Phillips, P. C. B., Shi, S., and Yu, J. (2015). Testing for multiple bubbles: Historical episodes of exuberance and collapse in the S&P 500. *International Economic Review*, 56(4):1043–1078.
- Phillips, P. C. B., Wu, Y., and Yu, J. (2011). Explosive behavior in the 1990s NASDAQ: When did exuberance escalate asset values? *International Economic Review*, 52(1):201–225.
- Quemin, S. (2020). Using supply-side policies to raise ambition: the case of the EU ETS and the 2021 review. Centre for Climate Change Economics and Policy Working Paper 366/Grantham Research Institute on Climate Change and the Environment Working Paper 335.
- Quemin, S. and Trotignon, R. (2019). Emissions Trading with Rolling Horizons. Centre for Climate Change Economics and Policy Working Paper 348/Grantham Research Institute on Climate Change and the Environment Working Paper 316.
- Rickels, W., Görlich, D., and Peterson, S. (2014). Explaining European Emission Allowance Price Dynamics: Evidence from Phase II. *German Economic Review*, 16(2):181–202.

- Robinson, P. M. (1989). *Nonparametric estimation of time varying parameters In: Hackl, P. (Ed.), Statistics, Analysis and Forecasting of Economic Structural Change*. Springer, Berlin.
- Said, S. E. and Dickey, D. A. (1984). Testing for unit roots in autoregressive-moving average models of unknown order. *Biometrika*, 71(3):599–607.
- Salant, S. W. (2016). What ails the European Union’s emissions trading system? *Journal of Environmental Economics and Management*, 80:6–19.
- Salant, S. W. and Henderson, D. W. (1978). Market Anticipations of Government Policies and the Price of Gold. *Journal of Political Economy*, 86(4):627–648.
- Samorodnitsky, G. and Taqqu, M. S. (1994). Stable non-gaussian random processes. *Chapman & Hall, New York*, 29.
- Santos, C., Hendry, D., and Johansen, S. (2008). Automatic selection of indicators in a fully saturated regression. *Computational Statistics*, 23:317–335.
- Sharma, S. and Escobari, D. (2018). Identifying price bubble periods in the energy sector. *Energy Economics*, 69:418–429.
- Shi, S. (2017). Speculative Bubbles or Market Fundamentals ? An Investigation of US Regional Housing Markets. *Economic Modelling*, 66:101–111.
- Thomson Reuters (2019). Carbon Market Survey 2018. Technical report.
- Voisin, E. and Hecq, A. (2019). Forecasting bubbles with mixed causal-noncausal autoregressive models. *Maastricht University. MPRA 96350*.
- Wetttestad, J. and Jevnaker, T. (2019). Smokescreen Politics? Ratcheting Up EU Emissions Trading in 2017. *Review of Policy Research*, 36(5):635–659.
- ZEW (2019). Energiemarktbarometer. Technical report.

Zhou, Z. and Wu, W. B. (2010). Simultaneous inference of linear models with time varying coefficients. *Journal of the Royal Statistical Society Series B* 72, pages 513–531.

Appendix A Technical Appendix

A.1 Time-varying coefficient estimation and confidence intervals

Given an observed set of data on the response series $\{y_t\}$ and the regressors $\{\mathbf{x}_t\}$, the coefficient curves $\beta_j(\cdot)$ in (4.1) can be estimated via local linear nonparametric kernel estimation as in Cai (2007). Underlying this method is a first-order Taylor approximation of each $\beta_j(\cdot)$, for $j = 0, \dots, d$, at a fixed time point $\tau \in (0, 1)$. For t/n in a neighborhood of τ it holds that

$$\beta_j(t/n) \approx \beta_j(\tau) + \beta_j^{(1)}(\tau)(t/n - \tau). \quad (\text{A.1})$$

with $\beta_j^{(1)}(\cdot)$ denoting the first derivative of the coefficient function $\beta_j(\cdot)$. If we replace $\beta_{j,t} = \beta_j(t/n)$ by approximation (A.1) for every $j = 0, \dots, d$, equation (4.1) can now be rewritten as:

$$y_t \approx \boldsymbol{\beta}(\tau)' \mathbf{x}_t + \boldsymbol{\beta}^{(1)}(\tau)' \mathbf{x}_t(t/n - \tau) + z_t, \quad (\text{A.2})$$

where $\boldsymbol{\beta}^{(1)}(\tau) = (\beta_0^{(1)}(\tau), \dots, \beta_d^{(1)}(\tau))'$ denote the stacked first derivatives of trend functions evaluated at τ . The local linear estimator of this model is found by minimizing the following weighted sum of squares with respect to $\boldsymbol{\theta}$, where $\tilde{\mathbf{x}}_t(\tau) = (\mathbf{x}_t, \mathbf{x}_t(t/n - \tau))'$:

$$\hat{\boldsymbol{\theta}}(\tau) = \operatorname{argmin}_{\boldsymbol{\theta}} \sum_{s=1}^n \{y_s - \tilde{\mathbf{x}}_s(\tau)' \boldsymbol{\theta}\}^2 K\left(\frac{s/n - \tau}{h}\right), \quad (\text{A.3})$$

where $K(\cdot)$ is a kernel function and $h > 0$ is a bandwidth. As $n \rightarrow \infty$, the bandwidth is assumed to satisfy $h \rightarrow 0$ while $nh \rightarrow \infty$. The solution to this minimisation problem gives the estimator of $(d + 1)$ coefficient functions $\boldsymbol{\beta}(\cdot)$ as well as their $(d + 1)$ first derivatives $\boldsymbol{\beta}^{(1)}(\cdot)$. Let $\boldsymbol{\theta}(\tau) = (\boldsymbol{\beta}(\tau), \boldsymbol{\beta}^{(1)}(\tau))'$ denote the vector of stacked coefficient functions and first derivatives. Then, the estimator $\hat{\boldsymbol{\theta}}(\tau) = (\hat{\boldsymbol{\beta}}(\tau), \hat{\boldsymbol{\beta}}^{(1)}(\tau))'$ can be expressed as,

$$\hat{\boldsymbol{\theta}}(\tau) = \begin{pmatrix} \mathbf{S}_{n,0}(\tau) & \mathbf{S}'_{n,1}(\tau) \\ \mathbf{S}_{n,1}(\tau) & \mathbf{S}_{n,2}(\tau) \end{pmatrix}^{-1} \begin{pmatrix} \mathbf{T}_{n,0}(\tau) \\ \mathbf{T}_{n,1}(\tau) \end{pmatrix} \equiv \mathbf{S}_n^{-1}(\tau) \mathbf{T}_n(\tau), \quad (\text{A.4})$$

for $\tau \in (0, 1)$, where for $k = 0, 1, 2$,

$$\begin{aligned}\mathbf{S}_{n,k}(\tau) &= \frac{1}{nh} \sum_{t=1}^n \mathbf{x}_t \mathbf{x}_t' (t/n - \tau)^k K\left(\frac{t/n - \tau}{h}\right), \\ \mathbf{T}_{n,k}(\tau) &= \frac{1}{nh} \sum_{t=1}^n \mathbf{x}_t (t/n - \tau)^k K\left(\frac{t/n - \tau}{h}\right) y_t.\end{aligned}$$

This estimator can be seen as a weighted least squares estimator of a model of the form (A.2). The fitted values are thus obtained by

$$\hat{y}_t = \mathbf{x}_t' \hat{\boldsymbol{\beta}}(t/n), \quad (\text{A.5})$$

which shows that we are only interested in the estimates of the coefficient curves, the estimated first derivatives only serve as a by-product of the specific choice of estimator.

For the construction of confidence intervals, we rely on the autoregressive wild bootstrap as proposed in Friedrich et al. (2019) for a nonparametric trend model. We extend the method to model (4.1) by using the following bootstrap algorithm:

Algorithm 1 (Autoregressive Wild Bootstrap).

1. Estimate model (4.1) and form a residual series. This means, calculate

$$\hat{z}_t = y_t - \mathbf{x}_t' \tilde{\boldsymbol{\beta}}(t/n), \quad t = 1, \dots, n,$$

where the estimate $\tilde{\boldsymbol{\beta}}(t/n)$ is obtained by bandwidth $\tilde{h} > h$.

2. For $0 < \gamma < 1$, generate ν_1^*, \dots, ν_n^* as i.i.d. $\mathcal{N}(0, 1 - \gamma^2)$ and let $\xi_t^* = \gamma \xi_{t-1}^* + \nu_t^*$ for $t = 2, \dots, n$. Take $\xi_1^* \sim \mathcal{N}(0, 1)$ to ensure stationarity of $\{\xi_t^*\}$.
3. Calculate the bootstrap errors z_t^* as $z_t^* = \xi_t^* \hat{z}_t$ and generate the bootstrap observations by $y_t^* = \mathbf{x}_t' \tilde{\boldsymbol{\beta}}(t/n) + z_t^*$ for $t = 1, \dots, n$, where $\tilde{\boldsymbol{\beta}}(t/n)$ is the same estimate as in the first step.
4. Obtain the bootstrap estimator $\hat{\boldsymbol{\beta}}^*(\cdot)$ as defined in (A.4) using the bootstrap series $\{y_t^*\}$, with the same bandwidth h as used for the original estimate $\hat{\boldsymbol{\beta}}(\cdot)$.

5. Repeat Steps 2 to 4 B times, and let

$$\hat{q}_{j,\alpha}(\tau) = \inf \left\{ u \in \mathbb{R} : \mathbb{P}^* \left[\hat{\beta}_j^*(\tau) - \tilde{\beta}_j(\tau) \leq u \right] \geq \alpha \right\}$$

denote, for all $j = 0, \dots, d$, the α -quantile of the B centered bootstrap statistics $\hat{\beta}_j^*(\tau) - \tilde{\beta}_j(\tau)$. These bootstrap quantiles are then used to construct confidence bands as described below.

Note that in Step 1 of the above algorithm, a different bandwidth is used to perform the nonparametric estimation. We follow the practical implementation in Friedrich et al. (2019) using $\tilde{h} = 0.5h^{5/9}$. Compared to the original bandwidth h , this bandwidth is larger and produces an oversmoothed estimate as starting point for the bootstrap procedure. The reason for this is the presence of an asymptotic bias whenever local polynomial estimation is applied. The bias contains the second derivatives of the coefficient functions, which can only be consistently estimated using a larger bandwidth \tilde{h} .

To construct confidence intervals for every $\beta_j(\cdot)$ for a confidence level of $1 - \alpha$, we use the quantiles as obtained in Step 5 of the bootstrap procedure. The confidence intervals are given by

$$I_{j,\alpha}^*(\tau) = \left[\hat{\beta}_j(\tau) - \hat{q}_{j,1-\alpha/2}, \hat{\beta}_j(\tau) - \hat{q}_{j,\alpha/2} \right]. \quad (\text{A.6})$$

A.2 Testing procedure of the bubble detection test

Phillips et al. (2015) propose a recursive procedure which tests for the presence of explosive episodes and locates the starting (and potentially termination) point of such episodes. The pricing equation underlying the approach can be written as

$$EUA_t = \sum_{i=0}^{\infty} \left(\frac{1}{1+r_f} \right)^i \mathbb{E}_t(U_{t+i}) + B_t, \quad (\text{A.7})$$

where r_f is the risk-free interest rate, U_t represents the underlying fundamentals and \mathbb{E}_t is the conditional expectation given the information available at time t . B_t is the bubble component which is assumed to satisfy $\mathbb{E}_t(B_{t+1}) = (1+r_f)B_t$. When $B_t = 0$, there is no bubble present in the

price series and the degree of nonstationarity is determined by the fundamentals which are assumed to at most $I(1)$. In the presence of bubbles, when $B_t \neq 0$, the price series shows explosive behaviour.

Under the null hypothesis of the test, the bubble component is zero such that the price follows an $AR(1)$ process with a drift:

$$EUA_t = dT^{-\eta} + \theta EUA_{t-1} + \epsilon_t,$$

where T is the sample size, d is a constant and $\eta > \frac{1}{2}$ controls the magnitude of the drift component. Following Phillips et al. (2015), we consider the case where $d = \eta = \theta = 1$ such that the price process will be $I(1)$.

Since the fundamentals are unobservable, it might be challenging to estimate this component of equation (A.7). In an ideal application, a widely used (theoretical) model exists with which one can obtain the fundamental value of this asset. In this case, the model can be calibrated and parameters can be estimated based on past data. This is done, for instance, in Shi (2017). In our application, however, no such model exists. As we have seen in Section 4.2, the relationship between allowance prices and the abatement-related fundamental price drivers is not stable over time. Given these results, we do not explicitly model the fundamental component but look for simultaneous explosive periods in the most established price drivers. We do this by applying the testing procedure separately to these series as it is done, for instance, in Corbet et al. (2018).

As explained in Phillips et al. (2015), in practice, it can also be difficult to distinguish bubbles from periods of price run ups, caused for instance by temporary changes in the discount rate. The latter can mimic bubble behaviour and will therefore also be detected by the test, although it is not incorporated in the theoretical framework where a constant risk-free rate is used. This is why Phillips et al. (2015) stress the importance of specifying in advance a minimum duration for an episode to qualify as bubble.

In general, the recursive testing procedure detects periods of mildly explosive behaviour and market exuberance in time series, and it is able to identify the location. The test applies a series of right-tailed ADF tests on a (backward and) forward expanding sub-sample. The regression model on which the test is based is closely related to the standard ADF test regres-

sion. It is a version of the same regression, but on a particular window:

$$\Delta EUA_t = \alpha_{r_1, r_2} + \beta_{r_1, r_2} EUA_{t-1} + \sum_{i=1}^k \psi_{r_1, r_2}^i \Delta EUA_{t-i} + \epsilon_t, \quad (\text{A.8})$$

where r_1 denotes the start of the window and r_2 the end, both expressed as a fraction of the sample size T . The ADF test is concerned with the null hypothesis $H_0 : \beta_{r_1, r_2} = 0$ and the statistic will be denoted by $ADF_{r_1}^{r_2}$. The minimum window size is r_0 and the actual window size is $r_w = r_2 - r_1$.

The above regression is run multiple times on $[Tr_w]$ observations. The SADF test, which is the first version of the test, was introduced by Phillips et al. (2011). The regression is estimated on a forward expanding sample, starting at $r_1 = 0$, whose length increases such that r_w runs from r_0 to 1. The test is the supremum over all ADF statistics, hence named the SADF test. Formally, it can be written as

$$SADF(r_0) = \sup_{r_2 \in [r_0, 1]} ADF_0^{r_2}. \quad (\text{A.9})$$

More powerful in the case of multiple bubbles is a variant of the test called generalised SADF (GSADF) test in which not only the end point of the window is varied but also the starting point r_1 . Different windows are considered, for r_1 varying from 0 to $(r_2 - r_0)$. The test statistic is defined as

$$GSADF(r_0) = \sup_{\substack{r_2 \in [r_0, 1] \\ r_1 \in [0, r_2 - r_0]}} ADF_{r_1}^{r_2}. \quad (\text{A.10})$$

Both tests are right-sided tests. Hence, if the test statistic exceeds the critical value, there is evidence for the existence of an explosive period. It does not provide any indication how many such episodes there are and where they are located. To achieve this, Phillips et al. (2015) develop a date-stamping strategy based on a third version of the sup ADF statistic, the backward SADF statistic (BSADF). This test proceeds in a similar way than the SADF and GSADF tests with the main difference of being obtained for every point in the sample. Fix a point in the sample as the

end point of the window, r_2 , and vary the starting point from 0 to $(r_2 - r_0)$, then the BSADF test is obtained as

$$BSADF_{r_2}(r_0) = \sup_{r_1 \in [0, r_2 - r_0]} ADF_{r_1}^{r_2}. \quad (\text{A.11})$$

Applying this test to each point in the sample results in a sequence of test statistics. To draw conclusions, we need to compare this sequence to a corresponding sequence of critical values. Before we can identify explosive episodes, it is important to define the minimum duration of a period to qualify as evidence for explosive behaviour. If the test statistic lies above the critical value merely for a few observations, this does not provide sufficient evidence for the existence of a bubble. It is rather a short-lived blip, as Phillips et al. (2015) call it. They suggest to chose a minimum duration which is dependent on the sample, such as $L_T = \log T$. We can then identify explosive episodes if we find periods for which the BSADF statistic exceeds the critical values for at least L_T consecutive observations. In our case the minimum duration according to this formula would be 7 weeks. This procedure can also identify ongoing bubbles and serve as an early warning system.

We apply the GSADF test and the date-stamping procedure in the next section using a minimum window size that is set according to the rule suggested by Phillips et al. (2015), $r_0 = 0.01 + 1.8\sqrt{T}$, resulting in 40 observations. For the selection of the number of lags we use the Bayesian Information Criterion. We apply two versions of the test. Firstly, the test as originally proposed by Phillips et al. (2015). For this, we obtain critical values from 2000 replications of their Monte Carlo simulation exercise to mimic the finite sample distribution given our sample size. The results are obtained using the R package `MultipleBubbles` which is accompanying the Phillips et al. (2015) paper. Secondly, we use a sieve bootstrap approach as proposed in Pedersen and Montes Schütte (2017). This paper shows that the GSADF test is critically oversized in the presence of serial correlation which implies that there is a serious risk of finding false evidence for explosive behaviour. The bootstrap corrects the size but it comes at the cost of a lower power. This is why we use both tests as complementary and show that they come to the same conclusion making our results more robust.

A.3 Modelling bubbles with noncausal α -stable processes

The α -stable noncausal AR(1) was introduced and partially studied by Gouriéroux and Zakoïan (2017) as a candidate model to analyse and forecast bubbles in financial time series, and was studied in details by Fries (2018a). This bubble-generating process features several surprising properties sharply contrasting with classical linear time series models; we start by recalling some of these properties before presenting the results regarding the prediction of bubble crash odds.

A noncausal AR(1) process is defined as the strictly stationary solution of the stochastic recurrence equation:

$$X_t = \rho X_{t+1} + \varepsilon_t,$$

where $0 < |\rho| < 1$, and (ε_t) is an i.i.d. error sequence. The appellation *noncausal*, standard in the literature (as well as *anticipative*), refers to the fact that the stationary solution (X_t) of the above equation admits a moving average representation in terms of «future» error terms as $X_t = \sum_{k=0}^{+\infty} \rho^k \varepsilon_{t+k}$. The above process is said to be an α -stable noncausal AR(1) when $\varepsilon_t \stackrel{i.i.d.}{\sim} \mathcal{S}(\alpha, \beta, \sigma, \mu)$, where $\mathcal{S}(\alpha, \beta, \sigma, \mu)$ denotes the univariate α -stable distribution with tail index $\alpha \in (0, 2)$, asymmetry $\beta \in [-1, 1]$, scale $\sigma > 0$ and location $\mu \in \mathbb{R}$, i.e., the distribution with characteristic function:

$$\mathbb{E}[e^{iu\varepsilon_0}] = \exp \left\{ -\sigma^\alpha |u|^\alpha \left(1 - i\beta \operatorname{sign}(u) w(\alpha, u) \right) + iu\mu \right\},$$

with $w(\alpha, s) = \operatorname{tg} \left(\frac{\pi\alpha}{2} \right)$, if $\alpha \neq 1$, and $w(1, s) = -\frac{2}{\pi} \ln |s|$ otherwise, for $s \in \mathbb{R}$. The case $\alpha = 2$ (excluded in the noncausal bubble modelling framework) corresponds to the Gaussian distribution, which is a particular stable distribution – and the only instance which does not feature heavy-tails. For $\alpha < 2$, the case of interest for bubble modelling, α -stable distributions have power-law decaying tails:

$$\mathbb{P}(|\varepsilon_0| > x) \underset{x \rightarrow \infty}{=} O(|x|^{-\alpha}).$$




Figure 6: Sample path of an α -stable noncausal AR(1) process with parameters $\rho = 0.95$, $\alpha = 1.7$, $\beta = 1$, $\sigma = 0.1$, $\mu = 0.5$.

For more details on univariate and multivariate α -stable distributions we refer to the comprehensive monograph by Samorodnitsky and Taqqu (1994). Figure 6 illustrates a typical sample path of a noncausal AR(1) process featuring multiple bubbles.

Although seemingly defined in reverse time, (X_t) is actually a Markov process in the classical sense: $X_{t+h}|X_t, X_{t-1}, \dots \stackrel{d}{=} X_{t+h}|X_t$ (Cambanis and Fakhre-Zakeri, 1995). As a linear combination of α -stable errors, X_t is itself marginally α -stable; it is thus heavy-tailed and admits very few marginal moments: $\mathbb{E}[|X_t|^s] < +\infty$ only for $s < \alpha$. Despite featuring infinite marginal variance, the conditional distribution of future realisation of the process, say, X_{t+h} for some horizon $h \geq 1$, given the past history X_t, X_{t-1}, \dots admits finite moments at higher orders, namely:

$$\mathbb{E}[|X_{t+h}|^\gamma | X_t] < +\infty, \quad \text{for } \gamma < 2\alpha + 1.$$

Thus, the predictive distribution of the process always admits a well-defined conditional expectation, even though marginally $\mathbb{E}[|X_t|] = +\infty$ for $\alpha \leq 1$, and can admit up to finite conditional variance, skewness and kurtosis, provided the tail exponent α is close enough to 2. Gouriéroux and Zakoïan (2017) obtained closed form expressions for the conditional expectation and

variance in the special cases of symmetric ($\beta = 0$) and Cauchy-distributed ($\alpha = 1, \beta = 0$) errors, highlighting that (X_t) can feature GARCH effects, and even presents a stationary unit root under special parameterisation (Cauchy errors and $\rho > 0$) in the sense that $\mathbb{E}[X_{t+1}|X_t] = X_t$.

The stationary solution (X_t) admits an infinite forward-looking moving average representation in terms of the future errors as $X_t = \sum_{k=0}^{\infty} \rho^k \varepsilon_{t+k}$. By a change of variable in the moving average representation, one can rewrite X_t as $X_t = \sum_{\tau \in \mathbb{Z}} \rho^{\tau-t} \mathbb{1}_{\tau-t \geq 0} \varepsilon_{\tau}$, and it can be noticed that the trajectory of the process (X_t) is a linear combination of so-called baseline paths (Gouriéroux and Zakoïan, 2017; Fries and Zakoïan, 2019): deterministic functions of time $t \mapsto \rho^{\tau-t} \mathbb{1}_{\tau-t \geq 0}$, which are shifted in time by τ and scaled by the random error ε_{τ} . Due to the heavy-tails of the distribution, realisations of the errors ε_t 's often take extreme values. Intuitively, if ε_{τ_0} is extreme for a particular date τ_0 , then the trajectory of (X_t) can be locally approximated by an exponential trend culminating in a peak and ending in a crash down to pre-bubble levels: $X_t \approx \rho^{\tau_0-t} \mathbb{1}_{\tau_0 \geq t} \varepsilon_{\tau_0}$, for t in the vicinity of τ_0 .

This intuition is formalised and shown to hold in Fries (2018a,b), and provides a convenient analytical framework to formulate predictions of likely future paths for trajectories of noncausal processes. In particular, it allows a formal quantification of the crash odds at future horizons of bubbles generated by the noncausal AR(1). Fries (2018a) obtained functional forms for the conditional expectation, variance, skewness and kurtosis of X_{t+h} given $X_t = x$ for any admissible parameterisation and for the general class of *two-sided* infinite moving average processes of the form $X_t = \sum_{k \in \mathbb{Z}} a_k \varepsilon_{t+k}$. Although the conditional moments have a complex form in general, their expressions drastically simplify when one considers the important case of interest of explosive episodes. For the noncausal AR(1) with $\rho > 0$, it is shown that during explosive episodes (i.e., large conditioning values $X_t = x$) the conditional moments of X_{t+h} given $X_t = x$ become equivalent to that of a weighted Bernoulli random variable charging probability $\rho^{\alpha h}$ to the value $\rho^{-h}x$, and probability $1 - \rho^{\alpha h}$ to 0. In other words, conditionally on an explosive episode having reached the level $X_t = x$, either:

1. with probability $\rho^{\alpha h}$, the bubble will survive at least until horizon h

and follow an exponential trend of growth rate ρ^{-1} : from $X_t = x$, the trajectory will grow to $\rho^{-h}x$.

2. with probability $1 - \rho^{\alpha h}$, the bubble will crash at some intermediate future date before horizon h : from $X_t = x$, the trajectory will crash to X_{t+h} , with $X_{t+h}/X_t \approx 0$.

This asymptotic behaviour of $X_{t+h}|X_t$, initially deduced from the form of first four conditional moments, has been confirmed to hold in distribution (Fries, 2018b). Surprisingly, it can be noticed that for the α -stable non-causal AR(1), the survival probability of bubbles at any horizon h , $\rho^{\alpha h}$, does not depend on the conditioning level $X_t = x$. Since the noncausal AR(1) is a Markov process, this implies that the survival probabilities of bubbles generated by the model do not depend at all on their past history: the bubbles feature a memory-less, or non-aging, property characterised by a geometric survival probability function. They can thus be fully characterised by their so-called half-life, the duration $h_{0.5}$ such that their survival probability at horizon $h_{0.5}$ is 50%. For the noncausal AR(1) with autoregressive parameter ρ and tail index α , the half-life reads:

$$h_{0.5} = -\frac{\ln 2}{\alpha \ln \rho}.$$

From the latter quantity, the crash odds (CO) at any horizon $h \geq 1$ can be expressed as

$$CO(h) = 1 - (1/2)^{\frac{h}{h_{0.5}}},$$

while the expected duration (ED) of an explosive episode reads:

$$ED(h) = \frac{1}{1 - \rho^{\alpha}}.$$

Appendix B Robustness Analysis

In addition to the final results shown in the paper, we conducted a formal outlier detection as well as an extensive bandwidth selection procedure, both serve as a robustness check for our results. In addition, we present results using additional/alternative data series.

B.1 Outlier detection

As mentioned in Section 4.2 we look for outliers in the dependent variable as our methods are not designed to explain sudden jumps in allowance price returns. The results presented in the main text are the results after outliers have been removed. We now explain how we detect them and how the results look without the removal.

We apply the impulse indicator saturation (IIS) approach proposed in Santos et al. (2008). This approach includes a dummy variable at every possible time point and performs expanding and contracting multiple block searches to determine which dummy variables should be retained. It is applied here to the EUA return series with a nominal significance level of $\alpha = 0.005$.




Figure 7: Summary of impulse indicator saturation results. Figure produced with R package *gets*. Top: EUA returns (blue) and fit of dummy regression (red); middle: standardised residuals after regression on dummies; bottom: retained dummy variables

The IIS method applied to our dependent variable retains 10 dummy variables corresponding to 10 outliers located at observations 39, 50, 53, 167, 237, 246, 254, 257, 305 and 444. Corresponding to time points in October 2008, January and February 2009, June 2011, November 2012, January 2013, March and April 2013, March 2014 and December 2016. If

we control for the impact of our most important explanatory variables – coal, gas, oil prices and temperature data – three of these outliers can be explained (at observation 39, 53 and 167). Figure 7 depicts the results of this IIS application. With this approach we retain 7 outliers which we delete from the EUA price series and subsequently repeat our analysis. The content of Figure 3 was produced with the outlier-corrected data. In Figure 8 we repeat this figure but the raw data. Comparison of the two figures shows that there are no major differences. A noteworthy difference is that the confidence intervals get more narrow over some periods. In particular, this effect is visible for the gas, coal and oil price coefficient in the period from end 2012 to the beginning of 2014, where for the original data the confidence intervals experienced a widening which disappears with the removal of outliers. This does not change the results in a substantial way nor does it affect the conclusions drawn regarding the significance of coefficients.

B.2 Bandwidth selection

A crucial aspect of any local fitting method is the choice of the bandwidth parameter h . It simultaneously controls the amount of smoothing and the complexity of the estimation. A small value of the bandwidth stands, on one hand, for a small approximation error and a resulting small modelling bias. On the other hand, it means that only a few data points are included in the local neighborhood and therefore, the variance of the estimator is large. In addition, the estimated curve is less smooth and the model is more complex than with a larger bandwidth.

In contrast, a large value of the bandwidth can create a large modelling bias, but the model is less complex and the amount of smoothing will be large. For the extreme values $h \rightarrow 0$ and $h \rightarrow \infty$, the estimate coincides, respectively, with interpolation of the data points or a parametric linear regression estimate. Therefore, bandwidth selection controls the complexity of the model and in addition, there is a bias-variance tradeoff.

Comparing parametric and nonparametric fitting, the bandwidth parameter can be seen as an additional dimension. Parametric fitting is like nonparametric fitting, where the choice of the bandwidth parameter is con-

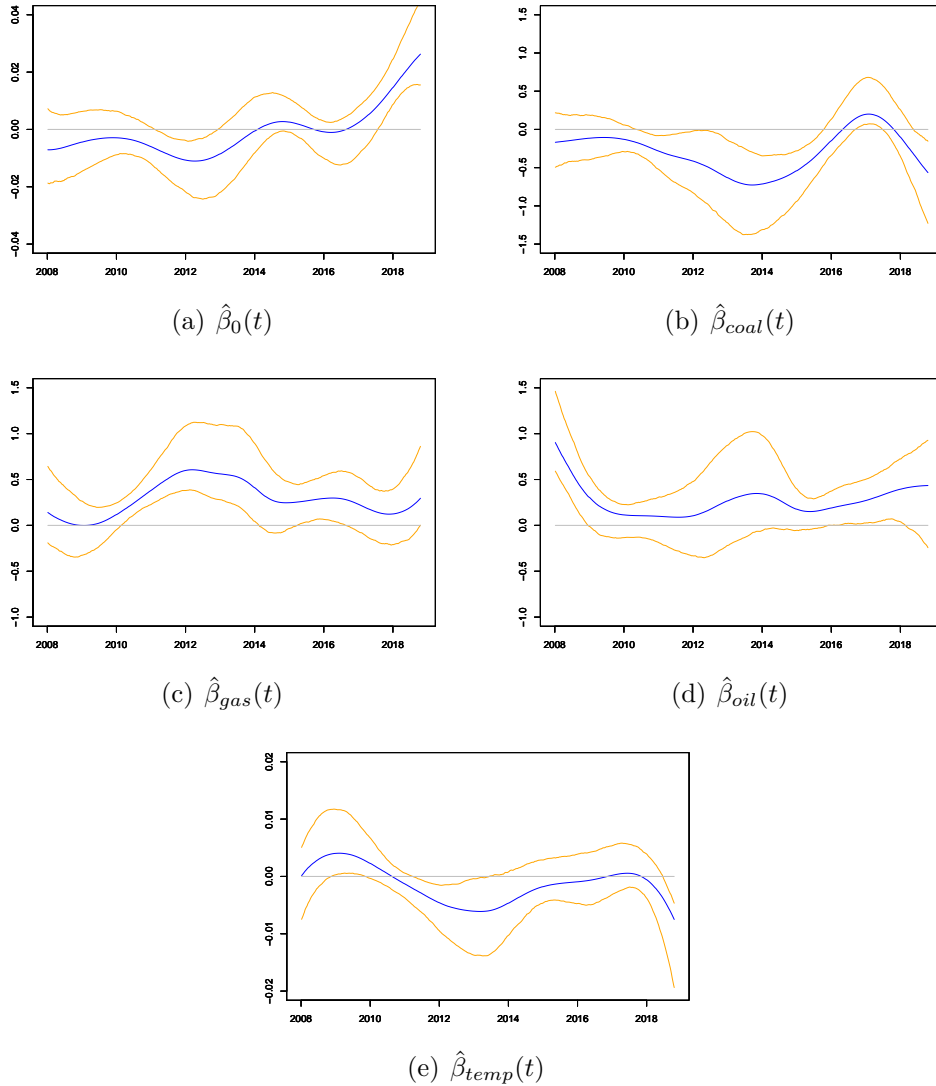


Figure 8: Nonparametrically estimated coefficient curves and 95% confidence intervals before the removal of outliers

stant and different families of models are considered. Taking a bandwidth of $h \rightarrow \infty$ is unquestioned in all situations of parametric modelling, while with nonparametrics, the bandwidth is seen as an additional parameter, which is carefully selected so that the estimation outcome fits the data well.

A theoretically optimal bandwidth can be obtained but it is infeasible for practical use as the expression depends on several unobservable quantities, e.g. the second derivative of $\beta(\cdot)$. We refer the interested reader to

Fan and Gijbels (1996) for more details. We focus on rules how to select bandwidths in practice.

Many data-driven methods for bandwidth selection are based on the principle of cross validation (CV). The basic idea is to find the value of the bandwidth that provides the best fit in terms of minimizing the sum of squared residuals without overfitting. Simply finding h that minimises the sum of squared residuals creates a problem of overfitting, since for $h \rightarrow 0$ a perfect fit is obtained in the limit. The result of this obviously problematic procedure would be the smallest considered bandwidth in all cases. The leave-one-out estimator provides a way to circumvent overfitting. The first step is to construct the leave-one-out estimator by leaving out the observation t that receives the highest weight in the local estimation. The second step in the least-squares CV approach is to look at the weighted average of the leave-one-out squared residuals and minimise them with respect to h .

Table 5: Bandwidth selection

	Bandwidth	Reference
Cross Validation	0.0874	e.g. Fan and Gijbels (1996)
Generalised CV	0.0800	Zhou and Wu (2010)
AIC	0.0866	Cai (2007)
LLO CV	0.0894	Chu and Marron (1991)
Average	0.0859	

Source: Own calculations using R.

Notes: Optimal bandwidth chosen by different methods. The leave- $(2l + 1)$ -out CV approach is applied with $l = 4$.

Cross validation, however, was originally designed for independent data and can therefore potentially lead to problems in time series applications. Chu and Marron (1991) show that cross validation systematically chooses bandwidths that are too small (too large) in the presence of positive (negative) autocorrelation. They propose a modification of the criterion called modified cross validation and show that it works well in time series applications. It follows the same general principle as CV, but it is based on a different estimator. Chu and Marron (1991) use a leave- $(2l + 1)$ -out version of the leave-one-out estimator.




Figure 9: The four different data-driven bandwidth selection criteria: (a) Cross validation, (b) Generalised cross validation, (c) Leave- $(2l + 1)$ -out CV, (d) Akaike-based bandwidth selection criterion

In addition to cross validation and modified cross validation, we follow the suggestions in Cai (2007) as well as Zhou and Wu (2010) and make use of two additional bandwidth selection methods. Cai (2007) considers an approach based on the Akaike information criterion (AIC) while Zhou and Wu (2010) use generalised cross validation (GCV) originally proposed by Craven and Wahba (1978).

The chosen bandwidths are presented in Table 5 while the criteria are plotted in Figure 9. All bandwidths are in a similar range, with the GCV criterion selecting the smallest and the leave- $(2l + 1)$ -out CV the largest bandwidth. In the main text we decide to use a bandwidth of 0.09 which corresponds to the average. We also run the analysis with a smaller and a larger bandwidth, $h = 0.07$ and $h = 0.11$. We observe that the results do not change much. For illustration, we plot the gas price coefficient obtained with the three different bandwidths in Figure 10. The results for the other estimates can be obtained from the authors upon request.

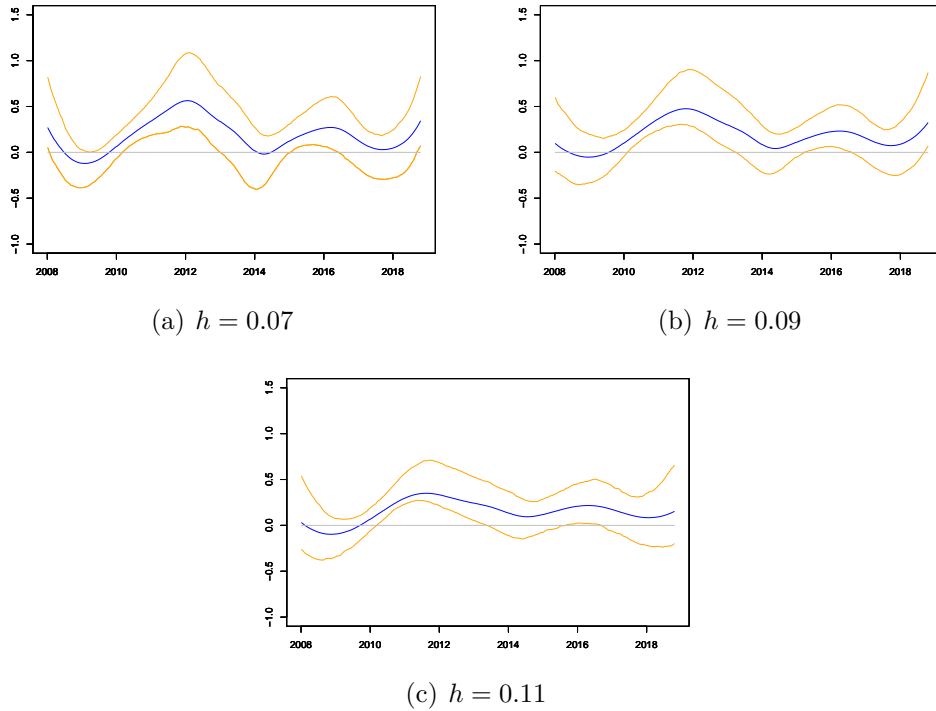


Figure 10: Three different estimates of the gas price coefficient, using $h = 0.07, 0.09, 0.11$

B.3 Additional explanatory variables

This section is two-fold. First, we replace the oil price as indicator of economic activity by the two stock indices in our data set – Euro STOXX Europe 50 and STOXX Europe 600. Second, we include several other variables (among them the fuel switching price) in the time-varying coefficient model which we did not include in the final results presented in the main paper. The effect of the fuel switching price is already captured in our model and the remaining additional variables did not show a period of significance. Excluding them from the regression does not alter the shape of the remaining coefficient curves.

Starting with the first point, we plot in Figure 11 the estimated parameter curves for the two stock indices. They were entered separately into the nonparametric regression, replacing the oil price. We give here only the coefficient of the replaced data series and not the entire set of regressors. Switching from the oil price to either one of the stock indices does not change the parameter estimates for the remaining regressors. Their

coefficient curves are almost identical to Figure 3. For the sake of brevity we do not plot them again. They can be obtained from the authors upon request.

Figure 11 reveals that the two stock indices produce very similar coefficient estimates. Given the shape of the two series (cf. Figure 2) this does not come as a surprise, as both indices show a comparable development over our sample period. Comparing the shape of the two curves in Figure 11 with the oil price coefficient plotted in Figure 3(d) shows that there are also some similarities. All three parameter estimates are positive and show some periods of significance which are longest for the oil price coefficient. While the oil price coefficient becomes significant at the end of 2014 and stays significant until the end of the sample (disregarding the boundary effect), the stock indices have a significant coefficient only from the end of 2014 to the beginning of 2015 and the STOXX 600 index in 2012 for a very short period. The period of significance is therefore overall shorter than for the oil price. Adding the two indices in addition to the oil price

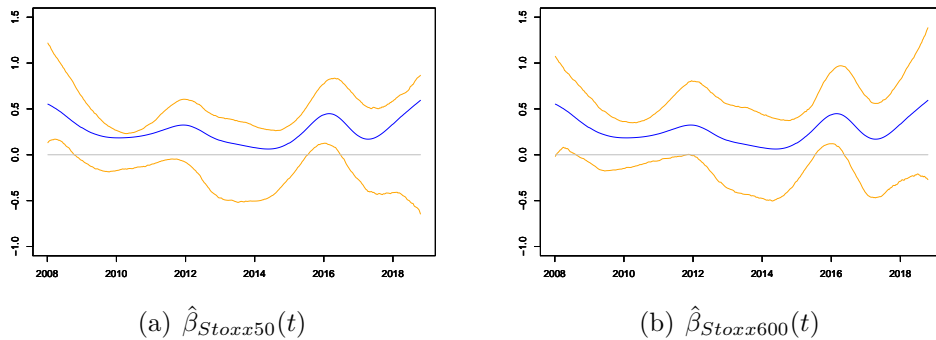


Figure 11: Nonparametrically estimated coefficient curves and 95% confidence intervals for the two stock indices

to the linear regression yields the results presented in Table 6. Neither of the indices has a significant effect on allowance prices when the oil price is already accounted for.

In the second and final part of this section, we add two additional explanatory variables to our nonparametric regression and we replace coal and gas prices by the fuel switching price. We obtain data on energy supply from hydro power in Norway from the Norwegian Water Resources

Table 6: Linear Regression Results

	(d)			(e)		
	$\hat{\beta}_j$	se_{NW}	p -value	$\hat{\beta}_j$	se_{NW}	p -value
Coal	-0.143	0.094	0.127	-0.124	0.109	0.255
Gas	0.176	0.078	0.024	0.213	0.074	0.004
Oil	0.206	0.068	0.003	0.175	0.088	0.048
Stoxx 50	–	–	–	0.07	0.158	0.671
Stoxx 600	–	–	–	–	–	–

	(f)		
	$\hat{\beta}_j$	se_{NW}	p -value
Coal	-0.145	0.097	0.136
Gas	0.177	0.077	0.022
Oil	0.178	0.086	0.040
Stoxx 50	–	–	–
Stoxx 600	0.112	0.165	0.498

Source: Own calculations using R.

Notes: Extension of Table 2. Results obtained using OLS estimation. The dependent variable is the return on EUAs and the set of (stationary) regressors changes in each specification. The p -values are based on Newey-West standard errors.

and Energy Directorate²¹ as well as data on electricity generation from wind for Germany obtained from the database of the European Network of Transmission System Operators for Electricity (ENTSO-E). Both variables should have a negative effect on allowance prices, as generation from renewable energy sources reduces emissions and therefore the demand for allowances. The hydro power data are weekly data which contain a strong seasonal component which is removed with the help of Fourier terms. This approach is also applied to the temperature data and it is explained in Section B.4 below. The wind generation data is only available until the end of May 2018 which reduces our sample size for this regression exercise compared to the main analysis. The sample comprises now 517 instead of 538 weekly observations. Added to the nonparametric regression, both new regressors have a coefficient estimate which is extremely low in mag-

²¹Data retrieved from <https://www.nve.no> on 21.01.2019.

nitude for the whole sample. Both estimated coefficient curves are plotted in Figure 12. From panel (a) we see that hydro power is significant over a very brief period in 2016. Panel (b) shows a period of significance for wind at the beginning of the sample until 2010. Although this period is quite long, given the small magnitude of the coefficient (in the order of 10^{-5}), we consider this effect as negligible.




Figure 12: Nonparametrically estimated coefficient curves and 95% confidence intervals for wind and hydro power generation

We calculate the fuel switching price for the switch from coal to gas for electricity generation. The switching price can be obtained from the coal and gas price series together with some additional values: the efficiency of coal and gas plants as well as the respective GHG emission factors. The emission factors are obtained from Juhrich (2016). The switching price is defined as follows

$$switch_t = \frac{\eta_{coal}p_{gas} - \eta_{gas}p_{coal}}{\eta_{gas}f_{coal} - \eta_{coal}f_{gas}}, \quad (B.1)$$

where η_i , f_i and p_i are the plant efficiency, emission factor and fuel prices for $i = coal, gas$, respectively. For details on this we refer to the review by Delarue et al. (2008). Figure 13 plots the resulting switching price as well as the estimated coefficient curve when replacing the gas and coal price series by the switching price. The switching price displayed in Panel (a) is obtained using our coal and gas price data as well as $\eta_{gas} = 0.47$, $\eta_{coal} = 0.36$, $f_{gas} = 0.202$ and $f_{coal} = 0.338$. Note that the gas and the coal price must both be denoted in EUR/MWh. This is why we need to divide our coal price by the conversion factor 8.14.




Figure 13: Fuel switching price (Panel (a)) and nonparametrically estimated coefficient curves and 95% confidence intervals for the fuel switching price (Panel (b))

The general development of the switching price is similar to the gas price as plotted in Figure 2(b). There is a visible upward trend at the end of the sample, which could again be seen as a potential cause of the explosive period in allowance prices. In Section 5, we already excluded the coal and the gas price as drivers of this behaviour. Therefore, we do not expect the results for the switching price to be different. For completeness, we obtain the GSADF statistic, which lies with 0.9647 below the critical values indicating no evidence for explosive periods in the switching price.

From (B.1) we can see that the switching price is a linear combination of the gas and the coal price. Hence, we include it in place of the two price series in our regression. The result is presented in Panel (b) of Figure 13. The shape of the estimated coefficient curve closely resembles the coefficient of the gas price which is not surprising given the stronger effect of the gas price on allowance prices compared to the coal price. The estimated coefficient is, however, smaller and the second period of significance vanishes. Including the gas and the coal price separately and not as a linear combination is less restrictive and we therefore do not consider the switching price in our main model.

B.4 Temperature data

The temperature data were obtained from the European Climate Assessment & Dataset (ECA&D) which provides surface air temperature for 199 measurement stations in Europe. It is provided by The Royal Netherlands

Meteorological Institute (KNMI). We refer to Klein Tank et al. (2002) for more details on the temperature series and measurement stations.


We obtain daily mean temperature series for seven cities, located in seven different countries. They are spread out over Europe: Berlin, Budapest, De Bilt, Dublin, Lyon, Madrid and Stockholm. All series do not contain any missing observations and all were updated until the end of 2018. We take the average over the cities as our temperature series. It is displayed in Figure 14(a). We aggregate the data to weekly means in order to match our sample frequency. In addition, we remove seasonality by fitting a Fourier regression and subsequently, working with the residual series from this regression. This removes the seasonal component with the help of sine and cosine functions. More specifically, we fit the following regression

$$Temp_t = \alpha_1 \cos(2\pi t) + \alpha_2 \sin(2\pi t) + \epsilon_t \quad (\text{B.2})$$


and continue to work with the residuals from this regression which are plotted in Figure 14(c). The remaining part (b) of Figure 14 plots the daily mean temperature (gray circles) together with the fitted Fourier terms (blue). We can see that the seasonal component is well captured by this method.

B.5 Additional BSADF results


Figure 15 presents the results from the date-stamping of explosive periods using the BSADF test. Compared to Figure 4 which is the bootstrap version, Figure 15 plots the series of BSADF test statistics (blue) and the critical value series (orange) obtained by Monte Carlo simulation. Panel (a) presents results from an application on the allowance price series and panel (b) on the oil price series. For both series, the conclusions are the same as presented in the main paper. The only difference lies in the starting point of the explosive period in allowance prices; while it was March 2018 in Figure 4 it is February in Figure 15. This minor difference is not surprising given the fact that the bootstrapped critical values are larger due to the size correction.



(a) Daily mean temperature




(b) Fourier fit




(c) Seasonally adjusted

Figure 14: Temperature data before and after the removal of the seasonal component.



(a) EUA prices



(b) Oil prices

Figure 15: Results from the BSADF test applied to the EUA prices (a) and the oil prices (b). Each panel shows the critical value series (orange) and the test statistics (blue)