Norges Bank

1 | 2024

Working Paper

The Drivers of Emission Reductions in the European Carbon Market

Norges Bank Research

Authors:
Hilde C. Bjørnland
Jamie L. Cross
Felix Kapfhammer

Keywords
Climate policy, Carbon pricing, emission trading system, cap and trade, demand and supply

Working papers fra Norges Bank, fra 1992/1 til 2009/2 kan bestilles på e-post: servicesenter@norges-bank.no

Fra 1999 og senere er publikasjonene tilgjengelige på www.norges-bank.no

Working papers inneholder forskningsarbeider og utredninger som vanligvis ikke har fått sin endelige form. Hensikten er blant annet at forfatteren kan motta kommentarer fra kolleger og andre interesserte. Synspunkter og konklusjoner i arbeidene står for forfatternes regning.

Working papers from Norges Bank, from 1992/1 to 2009/2 can be ordered by e-mail: servicesenter@norges-bank.no

Working papers from 1999 onwards are available on www.norges-bank.no

Norges Bank's working papers present research projects and reports (not usually in their final form) and are intended inter alia to enable the author to benefit from the comments of colleagues and other interested parties. Views and conclusions expressed in working papers are the responsibility of the authors alone.

ISSN 1502-8143 (online)

ISBN 978-82-8379-308-6 (online)

The Drivers of Emission Reductions in the European Carbon Market*

Hilde C. Bjørnland[†] Jamie L. Cross[‡§] Felix Kapfhammer[§]

January 29, 2024

Abstract

This paper studies the drivers of emission reductions in the carbon market of the European Union Emission Trading System (EU ETS) since its inception in 2005. We introduce a novel empirical framework that facilitates the joint identification of simultaneous demand and supply shocks underlying the European carbon market. We find that emission supply restrictions of the EU ETS were the dominant driver of emissions reductions, reducing emissions by 46%. However we also find that two opposing emission demand factors also played an important role. Demand from industrial economic activity increased emissions by 15%, while other demand-side factors, primarily reflecting the transition to low-carbon economies, reduced emissions by 21%.

JEL-codes: Q41, Q54, Q58

Keywords: climate policy, carbon pricing, emission trading system, cap and trade, demand and supply

*This Working Paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank. We are grateful for the comments and suggestions provided by Andrés Bellofatto, Dario Caldara, Begoña Dominguez, Lana Friesen, Francesco Furlanetto, Bård Harstad, Andrea La Nauze, Warwick McKibbin, Alicia Rambaldi, Francesco Ravazzolo, an anonymous referee for our Norges Bank Working Paper, the participants of the seminars at Luiss University, the University of Queensland, the Oslo Environmental and Climate Policy group, BI Norwegian Business School, and participants at the Climate Workshop at Monash Business School, the CCGE Workshop on Climate Change at Lancaster Management School and the 16th International Conference on Computational and Financial Econometrics in London. This paper is part of the research activities at the Centre for Applied Macroeconomics and Commodity Prices (CAMP) at the BI Norwegian Business School.

[†]CAMP, BI Norwegian Business School and Norges Bank.

[‡]Melbourne Business School, University of Melbourne.

[§]CAMP, BI Norwegian Business School.

1 Introduction

The European Union Emission Trading System (EU ETS) was established in January 2005, to reduce greenhouse gas (GHG) emissions within the region. At its inception, the EU ETS was the world's first carbon market and covered 4.4% of global GHG emissions (World Bank, 2022). In essence, the EU ETS is a cap and trade system that primarily requires the emission-intensive industries to acquire tradeable permits for each emitted ton of GHGs. These permits, known as European Union Allowances (EUAs), are either auctioned or distributed for free to the industry and can be traded in the open market. Since the introduction of the EU ETS, emissions have decreased by just over 50%. This raises the question of what has caused these emission reductions. Were they driven by the EU ETS regulation, or by other factors, such as industrial economic activity and shifting preferences towards low-carbon products? Answers to these questions are not only relevant for the design of the ETS, but also for market participants, private investors, and policymakers, who are involved in international policy debates.

The standard approach in empirical studies that quantify the causes of these emission reductions is to model demand and supply-side factors in isolation. On the one hand, numerous scholars have investigated the effects of the EU ETS supply-side restrictions on emissions (Ellerman and Buchner, 2008; Ellerman and McGuinness, 2008; Anderson and Di Maria, 2011; Abrell et al., 2011; Declercq et al., 2011; Gloaguen and Alberola, 2013; Petrick and Wagner, 2014; Wagner et al., 2014; Bel and Joseph, 2015; Jaraite-Kažukauske and Di Maria, 2016; Dechezleprêtre et al., 2023; Bayer and Aklin, 2020; Best et al., 2020; Colmer et al., 2022; Känzig, 2023). On the other hand, only Declercq et al. (2011) and Bel and Joseph (2015) have examined how demand-side factors affect emissions. Both papers investigate how the reduced economic activity during the 2008 Global Financial Crisis (GFC) impacted emissions.

In this paper, we present a unified framework for jointly identifying simultaneous supply and demand shocks in the EU ETS emissions market. Our framework quantifies the impact of both demand and supply factors on emission prices and quantities, and assesses their relative contributions to emissions reductions. To be best of our knowledge, this is the first paper that investigates demand factors other than economic activity, such as decreased demand for emissions resulting from the transition to a low-carbon economy. This is accomplished through two steps.

In the first instance, we present evidence indicating that the EUA cap in the emissions market, as imposed by the EU ETS, has often been non-binding in the short- and medium-term. Several factors contribute to this situation. Most notably, allowances not used for emitting GHGs can be saved, utilized, or (re)sold in the future. Consequently, the quantity of acquired yet unused EUAs steadily increased from zero at the inception of the cap-and-trade system to a peak of nearly two billion in 2013. This cumulative surplus suggests that the effective quantity of EUAs available in the carbon market exceeds the cap set by EU ETS policymakers for a given year. Therefore, the effective supply schedule in the emissions market is not perfectly inelastic, as often assumed; instead, the market typically operates on the upward-sloping portion of the supply curve. Recognizing that the supply of EUAs is price-sensitive has significant implications, as it implies that short- and medium-term emission reductions can result from shifts in both emission supply and emission demand within the carbon market.

In a second instance, we address this question by proposing a structural vector autoregressive (SVAR) model of the EU ETS carbon market that facilitates the joint identification of the simultaneous forces of demand and supply underlying the market for emissions. By accounting for possible reverse causality between the price and quantity of emissions and industrial economic activity, our framework also reflects the views of policymakers when deciding whether to revise EU ETS given the current state of the physical and economic climates. To address concerns relating to imperfect identifying information about the true values of the underlying structural (semi-)elasticities, we formally represent our identifying assumptions using Bayesian prior distributions. This strategy builds on the state-of-the-art framework for estimating SVAR models with the Bayesian methods of Baumeister and Hamilton (2015), and subsequent work that argues for the merits of this explicit approach to model identification over traditional methods (Baumeister and Hamilton, 2020, 2022a). It also avoids common errors when estimating elasticities using SVAR models (Baumeister and Hamilton, 2022b), thereby facilitating reliable estimates of the short-run (semi-)elasticities of the supply and demand curves in the market for emissions, and quantifying the drivers of emission reductions since the EU ETS's inception.

The model framework allows us to provide the first historical account of the relative contributions of demand and supply side factors in driving emissions covered by EU ETS.

Our main result is that emission supply restrictions set out by the EU ETS framework were the most important factor, reducing emissions by 46% between the introduction of the carbon market in January 2005 and December 2021. However, our results reveal that two distinct emission demand factors have also played a key role. In particular, demand from higher industrial economic activity increased emissions by 15% in the same period. However, this increase was more than offset by an emission reduction of 21% driven by other demand factors, which, we argue, primarily reflect the demand to transition to a low-carbon economy facilitated by firms reducing their dependence on fossil fuels due to behavioural changes and climate policies other than the EU ETS. Finally, we show that the results are robust to different model specifications, and that our preferred model framework, separating between supply, transition demand and industrial activity demand, gives more realistic results than a model aggregating the two demand shocks into one.

The remaining paper is structured as follows. Section 2 introduces the data used in our empirical analysis, provides a historical overview of the EU ETS, and describes the prevailing emission supply and demand forces. Section 3 contains the empirical framework used to model these simultaneous market forces. Section 4 presents and discusses the results, and section 5 concludes.

2 The EU ETS

The EU ETS was introduced in 2005 with the objective of reducing GHGs emissions of carbon-intensive industries. It is fundamentally a cap-and-trade system that can be viewed as operating at three tranches.[1]

At the highest level, the European Commission sets a total cap on the quantity of specific greenhouse gases that can be emitted by the economic sectors and installations covered by the system. Key sectors include electricity and heat generation, petroleum refining and the production of metals, cement, paper, and bulk chemicals. Some sectors that may be important causes of CO2 emissions, such as agriculture and transportation (apart from inner-European aviation), are currently not covered by the EU ETS. There are also no country-specific caps under the EU ETS. This is because such caps may create an incentive for firms to relocate their activities to countries with less stringent emission

[1] For a discussion of cap and trade system versus carbon taxes, see for instance Metcalf (2021), or see Metcalf and Stock (2023) for an analysis of the effect of carbon taxes.

rules; a phenomenon referred to as carbon leakage. The overall cap is decreased over time so that total emissions from these sectors decrease. Since its inception, the cap and trade system has covered between 38% and 47% of the total carbon emissions in the Eurozone.

After establishing an overall cap on the quantity of emissions, the Commission then allocates the rights to these emissions into tradeable permits, known as European Union Allowances (EUAs). Each EUA provides the holder with the right to emit one tonne of CO2 or the equivalent amount of other greenhouse gases, nitrous oxide and perfluorocarbons. The EUAs are then distributed to firms either freely, or through purchase in the primary auction market. All companies are eligible to partake in the auctions. After this process the permits become tradeable in the secondary market. This trading can be done either directly between buyer and seller, or via intermediaries.

The concluding tranche of the EU ETS relates to reporting and compliance. In short, firms are required to monitor and report their emissions annually. They must also provide a sufficient quantity of allowances to cover all of their emissions. If a company fails to provide sufficient allowances, then heavy fines are imposed.

In this paper we restrict our attention to the 19 Eurozone countries that have been part of the EU ETS from its introduction in January 2005 until the end of our sample in December 2021, and disregard emissions from inner-European aviation due to its special treatment and introduction at a later stage. Details on the various phases of the EU ETS are provided in sub-section 2.1. Moreover, since our objective is to determine the drivers of emissions reductions under the EU ETS, we emphasize that our focus is on the modeling of total verified industrial emissions under the current scope of the EU ETS, as opposed to the trading of individual EUAs in the primary of secondary markets. A broad discussion on the relevant aspects of these variables in relation to our study is provided in sub-section 2.2.

2.1 A Brief Chronology of the EU ETS

The real price of EUAs and the quantity of emissions produced in this period are shown in the top panel of Figure 1.[2] We observe that emissions have steadily declined over the

[2]Emissions covered by the EU ETS are only reported since its start in 2005. To initialize our model and maximize the sample size we backcast emissions to 2003. This is done through a counterfactual in which we compute the quantity of emissions that would have fallen under the EU ETS in 2003 and 2004 based on the share of total emissions covered by the EU ETS across the years 2005-2007. We also follow Känzig (2023) and temporally disaggregate the annual frequency emissions data to a monthly

sample period, with substantial drops during the 2008 Global Financial Crisis (GFC) and the more recent COVID-19 pandemic period. This suggests that demand-side factors play a key role in driving emissions over the sample period.

Figure 1 also highlights the four phases that the EU ETS has undergone since its inception (European Commission, 2022a). Phase 1 (2005-07) was a pilot phase that was primarily designed to establish a carbon market and an associated price for carbon. During this phase, only CO_2 emissions were covered by the EU ETS, while further GHGs were added in subsequent phases. In the first months after the introduction of the carbon market, emissions slightly declined and the carbon price reached its first peak of around 30€ per tonne of CO_2-equivalent in 2006. However, the market collapsed subsequently and the EUA price reverted back to zero since the quantity of supplied emission allowances outstripped demand (Ellerman and Buchner, 2008). Consequently, market participants saved the acquired but unused EUAs for future use or resale. The build-up of the allowance surplus can be seen in the bottom panel of Figure 1, where we plot the cumulative quantity of unused EUAs over the sample period. As an exemption from the rule, the unused EUAs of the pilot phase could not be transferred to Phase 2, which strongly contributed to the collapse of the carbon price.

Phase 2 (2008-12) can be viewed as a policy correction phase in which the oversupply issues of Phase 1 were addressed by tightening the cap on emissions and allowing again for EUAs to be saved for future use or resale across the phases. The adjustments lifted the carbon price back up to the 30€ price observed in 2006, and emissions strongly declined. In 2008, the total number of unused allowances was even negative. This can only happen if companies receive their allocation of allowances for the subsequent year before they have to submit EUAs covering their emissions of the past year. Hence, companies can borrow from future EUA allocations. Two defining episodes of this period were the GFC and the associated European Debt Crisis (EDC). Both events led to a substantial fall in the EUA price and emissions as the demand for emissions collapsed due to lower industrial production (as shown by the black line in the bottom panel of Figure 1). This collapse in economic activity led to a hoarding of EUAs from near zero at the end of 2008 to

frequency using the Chow and Lin (1976) method. Since no EUA price for emissions is readily available for the whole period, we proxy the price with spot prices and short-term futures, which tend to comove very closely. The nominal price was then deflated by the Euro area harmonised index of consumer prices (HICP). Finally, we use coverage-weighted measures where available to ensure consistency across variables. Coverage-weighted measures are also employed in Metcalf and Stock (2023) to examine the effects of European carbon taxes. Further details are provided in the Online Appendix.

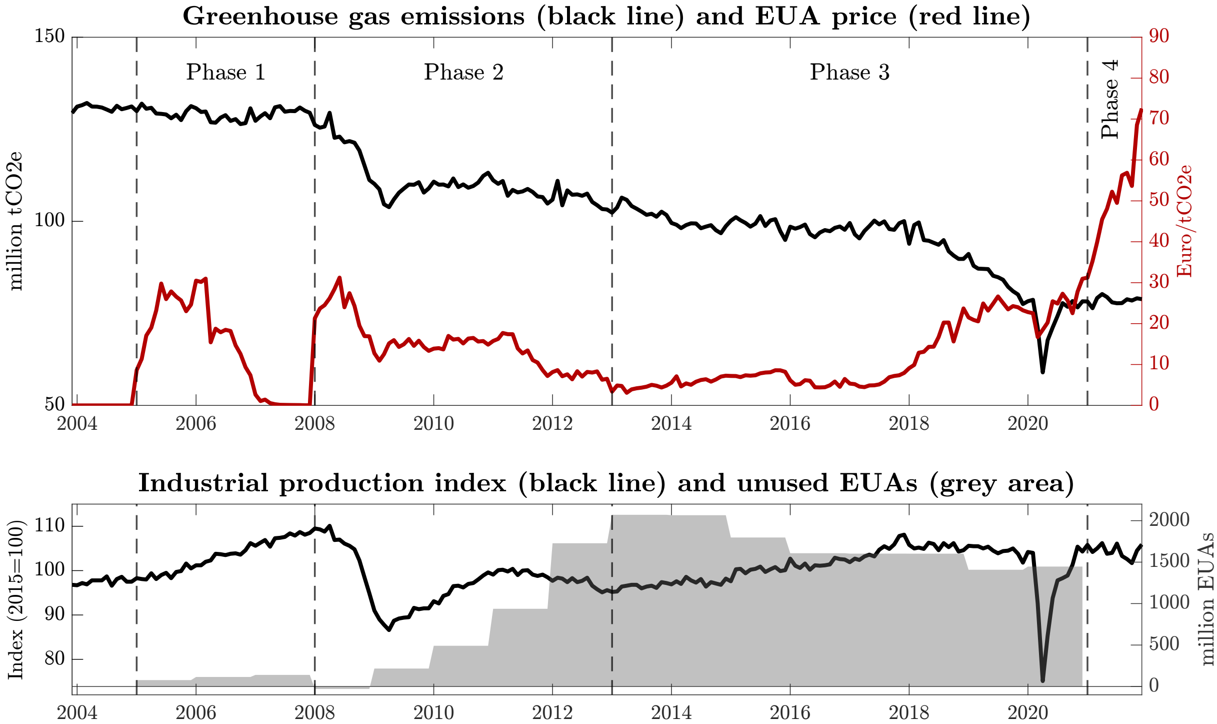

Figure 1: Data depicting the EU ETS.

Note: The upper panel shows GHG emissions in the Eurozone countries that are covered by the European emission trading system (black line, left axis) and the real EUA price (red line, right axis). The lower panel shows the industrial production index of the Eurozone countries (seasonally and calendar-adjusted with the base year 2015, left axis) and the total amount of unused EUAs in the carbon market between 2005 and 2020 (right axis). The phases are reflective of major policy changes in the EU ETS legislation (European Commission, 2022a).

around two billion by 2013. A further reason for the accumulation of allowances was the possibility for firms to offset emissions by buying a limited amount of *international credits* (certified removals of atmospheric GHGs) instead of using EUAs.

Phase 3 (2013-2020) covers the aftermath of the GFC and EDC, during which EU policymakers attempted to deal with the excessive surplus of unused EUAs. In this period, international credits must be exchanged for EUAs to avoid additional emission supply. The distribution of a large share of allowances was also postponed between 2014 and 2016 in an effort to reduce the *back-loading* of allowances and stabilise the carbon price. The retention resulted in a slight gradual increase in the price of allowances and a reduction in the quantity of emissions and unused EUAs. Nevertheless, the carbon price remained low until 2017 despite the reduced supply of allowances. In 2019, the EU established the *Market Stability Reserve* to adjust the total amount of auctioned allowances based on pre-defined rules to increase the carbon price and provide more stable incentives for green investments (European Commission, 2022b).

Finally, Phase 4 (2021-2030) has, so far, been limited to the aftermath of the most recent pandemic period. As shown in Figure 1, the Covid-19 crisis at the beginning of 2020 led to a massive drop in economic activity, which reduced emissions enormously. However, while emissions collapsed by one-third, the EUA price dipped only slightly, presumably due to the option of saving unused allowances for the future. Perhaps most stark is the fact that the aftermath of the Covid-19 pandemic has seen the EUA price surge to a historical high of just over 70€, despite the quantity of emissions stabilising at its pre-pandemic level. The soaring carbon price might be explained by several factors, including a faster reduction of the EU ETS cap relative to the previous phases, an increased intervention by the Market Stability Reserve by preserving a greater quantity of EUAs (see Revision for Phase 4 by the EU Commission), an increased electricity generation by coal due to high natural gas prices, and an increased energy demand due to the economic recovery after the Covid-19 crisis.

This brief chronology of the European carbon market highlights four key reasons why its supplied quantity of emissions is not bound to the cap on allowances set by the EU ETS: First, EUAs that are not redeemed in a given year can be saved for future use or resale. Second, firms might receive the next period's allocation of EUAs before they must balance their emissions, allowing them to borrow EUAs from the future. Third, international credits were temporarily used in addition to EUAs. Fourth, the Market Stability Reserve allows policymakers to adjust the quantity of EUAs for price stabilization purposes.

2.2 The Market for Emissions: Supply and demand shocks

Conventional wisdom might suggest that the observed emission reductions under the EU ETS are entirely driven by the supply constraints set out by the ETS's cap on the quantity of EUAs. Indeed, if all EUAs were used in the period in which they were issued, then the quantity of emissions would equal the quantity of supplied EUAs. In such a case, the perfectly inelastic supply constraint set in the market for EUAs would map one-to-one into a perfectly inelastic supply schedule in the associated market for emissions. However, the brief chronology presented in the previous subsection provides strong evidence that, in practice, the market generally emits well under or above the cap set by the EU ETS. This wedge between the actual quantity of emissions and the cap of the EU ETS framework points towards flexibility in the supply of emissions. Specifically, the maximum possible quantity of emissions in a given period is not determined by the emissions cap, but instead driven by an *effective cap*, which consists of the aggregated quantity of allowances provided by the ETS's cap, the surplus of unused EUAs from previous periods, the Market Stability Reserve's intervention, the possibility of borrowing from future periods, and international credits. Furthermore, the previous supply factors are likely price-sensitive. For example, market participants will sell more of their stocked EUAs if the EUA price is high, and the Market Stability Reserve will retain more allowances if the EUA price is low. Thus, the supply schedule is not perfectly inelastic, suggesting the market operates along the upward-sloping portion of the supply curve. This implies that the market price and quantity of emissions are determined by both the supply and demand of emissions.

While emission supply in the European carbon market is solely governed by the setup of the EU ETS, emission demand depends on several factors. As previously indicated, the main contributor to emissions is industrial economic activity as the EU ETS covers primarily the carbon-intensive industry. If economic activity unexpectedly collapses as, for example, during the GFC, then emissions decline substantially (Declercq et al., 2011; Bel and Joseph, 2015). However, other demand-side factors could also lead to lower emissions even if industrial economic activity remained constant.

For instance, consumers, producers and investors might change their behaviour as they become increasingly aware of the detrimental effects of emissions. Investors might direct more capital to green investments, consumers might voluntarily pay a premium for low-carbon products and producers might voluntarily replace high-carbon products with low-carbon alternatives. Furthermore, governments might implement climate policies other than carbon pricing for the industrial sector to support the transition to a low-carbon economy, c.f. Baranek et al. (2021). Such climate policies include subsidies for the development and implementation of green technologies and the regulation and prohibition of certain fossil fuels or emission levels. Both behavioural changes and climate policies other than the EU ETS decrease the demand for emission, but do not necessarily impair industrial economic activity. On the contrary, investors reducing the capital costs for green investments, consumers paying voluntarily a green premium for products and subsidies for the development and implementation of green technology might even increase industrial economic activity.

With this in mind, we focus on two distinct demand factors that can shift emissions in the European carbon market: *industrial economic activity* and *transition demand* – encompassing behavioural changes, climate policies other than the EU ETS, and potentially further unaccounted factors. Next, we discuss the identification strategy to disentangle these two emission demand shocks and the emission supply shocks.

3 Empirical Framework

3.1 Modelling the European Emission Market

The discussion in Section 2 suggests that the EU ETS carbon market can be described with $n = 3$ monthly variables: the log-growth of verified emissions q_t, the first-difference of the EUA price p_t and the log-growth of the industrial production index y_t. The effective sample period begins with the start of the EU ETS in January 2005 and ends with the latest available data in December 2021. All data sources and details on transformations are provided in the Online Appendix. The structural relationships of demand and supply in the EU ETS carbon market are mapped with the following three equations:

$$q_t = \overset{+}{\alpha}_{12} p_t + \boldsymbol{b}_1' \boldsymbol{x}_{t-1} + u_{1t} \tag{1}$$

$$q_t = \overset{-}{\alpha}_{22} p_t + \overset{+}{\alpha}_{23} y_t + \boldsymbol{b}_2' \boldsymbol{x}_{t-1} + u_{2t} \tag{2}$$

$$y_t = \overset{-}{\alpha}_{32} p_t + \boldsymbol{b}_3' \boldsymbol{x}_{t-1} + u_{3t} \tag{3}$$

in which $\boldsymbol{x}_{t-1}$ contains an intercept and twelve lags of the three variables of interest.

Equation (1) represents the supply of emissions, which increases with the EUA price as indicated by the positive sign above the supply semi-price elasticity coefficient $\overset{+}{\alpha}_{12}$. If the supply curve is perfectly inelastic then $\alpha_{12} = 0$, while $\alpha_{12} > 0$ suggests an upward-sloping supply curve. As discussed in Section 2, the positive semi-price elasticity is a consequence of the mechanism through which an increase in the EUA price stimulates the sales of saved allowances, the provision of international credits and the release of more allowances through the Market Stability Reserve. The average reduction in the supply of emission allowances, that is, the decreasing cap set by the EU ETS, is primarily captured by the constant of this equation, which is embedded in $\boldsymbol{x}_{t-1}$. Any variation in emissions that is not covered by the described structural relationships or the lags is captured by the *emission supply shock* u_{1t}.

Equation (2) describes the demand for emissions, which is downward sloping, as indicated by the negative semi-price elasticity α_{22}^{-}. The negative sign captures the idea that a higher EUA price causes firms to reduce emission intensive production, which is reflected in a lower demand for emissions. The coefficients α_{12}^{+} and α_{22}^{-} allow emission demand and supply to react simultaneously on EUA price changes. To allow for the fact that increased industrial economic activity necessarily generates emission demand (see Section 2), we introduce the production elasticity α_{23}^{+} that is restricted to be positive. The structural error u_{2t} of equation (2) therefore captures shocks to emissions demand that are orthogonal to the EUA price and industrial economic activity. Later on, when discussing the results, we present evidence suggesting that this shock primarily captures variation in emissions caused by the replacement of fossil fuel-based energy by renewable energy, thereby reflecting the demand to transition to a low-carbon economy. Hence, we refer to u_{2t} as a *transition demand shock*.[3] The transition to a low-carbon economy is expected to occur rather slowly and steadily instead of abruptly. Therefore, much of the transition is likely determined by its average speed, which is captured by the constant of this equation that is embedded in $\boldsymbol{x}_{t-1}$.

Finally, equation (3) models industrial economic activity, which decreases with the EUA price due to increased input costs for the industry as indicated by the semi-price elasticity of industrial production α_{32}^{-}. Industrial activity does not depend contemporaneously on emissions as they are not relevant input factor for production. Still, changes in emissions can affect industrial production, but only indirectly through the carbon price. All unaccounted changes in equation (3) form the shock to industrial economic activity u_{3t}. This shock changes the demand for emissions in the context of our model, and therefore, we refer to it as *industrial activity demand shock*.

The model setup implies that the emission supply shock u_{1t} is the only source shifting the emission supply curve, whereas the emission demand curve can be shifted by two shocks, the transition demand shock u_{2t} and the industrial activity demand shock u_{3t}. All three shocks can change the quantity and price of emissions so long as α_{12}^{+}, α_{22}^{-} and α_{23}^{-} are not equal to zero.

[3] While this shock could also reflect other factors related to emission demand, such as a changes in consumer/firm behavior that are not related to the transition to a low-carbon economy, in Section 4 we provide strong evidence that it primarily reflects transition demand.

3.2 SVAR Representation

For estimation purposes, it is important to note that the system of equations (1)-(3) has the following structural vector autoregressive (SVAR) representation

$$\boldsymbol{A}\boldsymbol{y}_t = \boldsymbol{B}\boldsymbol{x}_{t-1} + \boldsymbol{u}_t, \qquad \text{with } \boldsymbol{u}_t \sim i.i.d.\ N(\boldsymbol{0}, \boldsymbol{D}), \tag{4}$$

where $\boldsymbol{y}_t = (q_t, p_t, y_t)'$, $\boldsymbol{A}$ is the structural impact matrix, $\boldsymbol{B}$ is the coefficient matrix on elements of $\boldsymbol{x}_{t-1}$, and $\boldsymbol{u}_t = (u_{1t}, u_{2t}, u_{3t})'$ is a vector of independent and identically distributed Gaussian shocks with mean zero and a diagonal covariance matrix $\boldsymbol{D}$. The structural relationships of $\boldsymbol{y}_t$ are summarised by

$$\boldsymbol{A} = \begin{bmatrix} 1 & -(\alpha_{12}^{+}) & 0 \\ 1 & -(\alpha_{22}^{-}) & -(\alpha_{23}^{+}) \\ 0 & -(\alpha_{32}^{-}) & 1 \end{bmatrix}. \tag{5}$$

It is well known that the contemporaneous impact of the structural shocks is also linked to the inverse of the impact matrix $\boldsymbol{H} = \boldsymbol{A}^{-1}$, which is determined by

$$\boldsymbol{H} = \frac{1}{\alpha_{12} - \alpha_{22} - \alpha_{23}\alpha_{32}} \begin{bmatrix} -\alpha_{22} - \alpha_{23}\alpha_{32} & \alpha_{12} & \alpha_{12}\alpha_{23} \\ -1 & 1 & \alpha_{23} \\ -\alpha_{32} & \alpha_{32} & \alpha_{12} - \alpha_{22} \end{bmatrix}.$$

Given the imposed sign restrictions on the four coefficients, α_{12}^{+}, α_{22}^{-}, α_{23}^{+} and α_{32}^{-}, we can infer the pre-determined signs of the elements of $\boldsymbol{H}$ to evaluate if they align with our theoretical reflections about the three shocks.

$$\begin{bmatrix} + & + & + \\ - & + & + \\ + & - & + \end{bmatrix} \tag{6}$$

The signs of matrix $\boldsymbol{H}$ indeed match the theoretical presumptions of how the supply, transition demand and industrial activity demand shocks (sorted by column) affect emissions, the EUA price and industrial production (sorted by row). For example, the first column shows that a positive emission supply shock - an unexpected provision of more EUAs on the market - increases emissions and decreases the EUA price, which, in turn, increases industrial production due to lower input costs for firms.

In choosing a VAR system to model the European carbon market, our approach relates to the recent paper by Känzig (2023). Two important dimensions differ, however. First, we investigate supply- and demand-side factors that drove the observed reduction in emissions, while Känzig (2023) examines the effect of the carbon price shock. Second, we propose a combination of sign and exclusion restrictions to identify both supply and demand shocks underlying the emissions market, while Känzig (2023) uses an instrument to identify the sole impact of a supply-side driven carbon price shock. The two studies can therefore be viewed as complementary since they address different questions but refer to the same carbon market.

3.3 Priors

The conventional approach to structural identification required scholars to treat some selected parameter values as known with absolute certainty while treating the remaining parameter values as being completely unknown. However, this approach eliminates the incorporation of foundational economic theory into the analysis. For instance, in most markets, it is typically "known" *a priori* (or at least strongly believed) that the supply curve is upward-sloping and the demand curve is downward-sloping. Here we follow a recent trend in the literature and make our identifying assumptions explicit by specifying them through Bayesian priors (Baumeister and Hamilton, 2015, 2020, 2022a). Importantly, this approach also avoids common errors when estimating elasticities using SVAR models (Baumeister and Hamilton, 2022b), thereby facilitating reliable estimates of the (semi-)elasticities of the supply and demand curves in the emissions market.

To the best of our knowledge, there are no microeconometric estimates of the (semi-) elasticities in (5). Hence, we express our identifying assumptions through a prior $p(\boldsymbol{A})$, which consists of truncated Student's t-distributions that are independently imposed on each element of $\boldsymbol{A}$. The sign restrictions are motivated by the aforementioned theoretical considerations in the emissions market: an upward-sloping supply curve ($\alpha_{12} \geq 0$), a downward-sloping demand curve ($\alpha_{22} \leq 0$), higher industrial economic activity increases emissions ($\alpha_{23} \geq 0$), and higher carbon prices depress industrial economic activity ($\alpha_{32} \leq 0$). Information about the shape of the distribution is expressed through the location, scale and degrees of freedom (df) parameters. To accommodate the conventional perspective of a perfectly inelastic supply curve, we take a conservative stance and locate mass of the prior $p(\alpha_{12})$ at zero while leaving sufficient probability mass between 0 and 1 to allow for a (non-perfectly) inelastic supply of emissions (shape of 0.2 and df of 3). Since emission demand may react contemporaneously to variations in the EUA price, we locate the prior $p(\alpha_{22})$ at -0.1 while allowing for higher or lower elasticity (shape of 0.2 and df of 3). This means that, on average, a 1€ increase in the EUA price leads to a 0.1 percentage point (pp) decrease in emissions demand on impact. A similar rationale is used for the remaining two elasticities. For the case of $p(\alpha_{23})$, we specify a unitary elastic relationship (location at 1.0) while allowing for possible deviations (scale of 0.3 and df of 3). This is motivated by the reflection that a 1% increase in industrial economic activity will likely lead to a similarly large increase in emissions. Lastly, we set the location of the prior for the semi-price elasticity of industrial production $p(\alpha_{32})$ to -0.1 with df of 0.2 and 3.

The priors for the remaining model parameters are in line with those used elsewhere in the literature on estimating VAR models with Bayesian methods (e.g., Baumeister and Hamilton, 2019; Aastveit et al., 2023). Specifically, independent Gamma priors are placed on the diagonal elements of the structural variance matrix $\boldsymbol{D}$, and independent Minnesota-type Normal priors on the rows $\boldsymbol{b}_i$ of the lagged coefficient matrix $\boldsymbol{B}$.

3.4 Historical Decomposition

As discussed in Section 2.2, emission reductions are largely driven both systematic policy changes in the EU ETS framework in addition to unsystematic shocks. To assess the drivers of emission reductions over time we therefore account for both systematic and unsystematic effects in the historical decomposition. To that end, note that reduced-form representation of the SVAR model in (4) is given by:

$$\begin{aligned} \boldsymbol{y}_t &= \boldsymbol{H}\boldsymbol{c} + \boldsymbol{\Phi}\boldsymbol{x}_{t-1} + \boldsymbol{H}\boldsymbol{u}_t, \\ &= \boldsymbol{H}\boldsymbol{c} + \boldsymbol{\phi}_1\boldsymbol{y}_{t-1} + ... + \boldsymbol{\phi}_2\boldsymbol{y}_{t-12} + \boldsymbol{H}\boldsymbol{u}_t, \end{aligned} \tag{7}$$

where $\boldsymbol{H} = \boldsymbol{A}^{-1}$, $\boldsymbol{\Phi} = \boldsymbol{A}^{-1}\boldsymbol{B}$ and $\boldsymbol{\phi}_i$, $i = 1, \cdots, 12$ denote the autoregressive matrices. The h-step-ahead impulse response function, $\tilde{\boldsymbol{\Theta}}_h$, is given by:

$$\tilde{\boldsymbol{\Theta}}_h = \boldsymbol{J}\boldsymbol{F}^h\boldsymbol{J}'\tilde{\boldsymbol{H}}, \tag{8}$$

where $\boldsymbol{F}$ is the companion form matrix, $\boldsymbol{J} = [\boldsymbol{I}_n \quad \boldsymbol{0}_{n\times n(p-1)}]$, and $\tilde{\boldsymbol{H}} = \boldsymbol{H}\boldsymbol{Z}^{-1}$ provides the initial impact of the negative emission supply shock and positive emissions demand shocks at $h = 0$. Given these definitions, the historical decomposition is given by:

$$\boldsymbol{y}_t = \underbrace{\sum_{j=0}^{t-1} \boldsymbol{J}\boldsymbol{F}^j\boldsymbol{J}'\boldsymbol{H}\boldsymbol{u}_{t-j}}_{\text{Contr. of shocks}} + \underbrace{\sum_{j=0}^{t-1} \boldsymbol{J}\boldsymbol{F}^j\boldsymbol{J}'\boldsymbol{H}\boldsymbol{c}}_{\text{Contr. of constant}} + \underbrace{\sum_{j=0}^{p} \boldsymbol{J}\boldsymbol{F}^{t+j}\boldsymbol{J}'\boldsymbol{y}_{0-j}}_{\text{Contr. of initial conditions}}, \tag{9}$$

where $\boldsymbol{y}_{0-j}$ denotes the initial conditions. The historical decomposition decomposes the total variation in the variables of interest into contributions from three distinct components: unsystematic structural shocks, systematic structural constants, and initial conditions. Since most studies focus on the effects of economic shocks, the historical decomposition is often used to only illustrate the contribution of the unsystematic structural shocks, with the latter two components being disregarded (see, e.g., Baumeister and Hamilton, 2019 or Känzig, 2023). However in our framework we include the intercepts since we are also interested in contributions of the structural constant $\boldsymbol{c}$ due to their econometric interpretations. For example, the constant c_1 reflects the scheduled decrease in emission supply over time due to the declining EU ETS cap, whereas c_2 could reflect trends in technological progress or behavioral change.

4 Results

The SVAR model is estimated using the state-of-the-art Metropolis-within-Gibbs Monte Carlo Markov Chain (MCMC) algorithm of Baumeister and Hamilton (2015). We use a total of one million draws from the posterior distribution. The first half of these draws is discarded as a burn-in period, while the second half is used for inference. In the following, we present our main results on structural parameter estimates, impulse response functions and a historical decomposition of the observed emission reductions under the EU ETS.[4]

[4]In the Online Appendix we provide additional results on convergence diagnostics, a historical decomposition of the EUA Price, and the structural shocks series. There we also provide robustness checks on alternative lag lengths and the impact of the Covid-19 period, as well as introducing a set of narrative restrictions during the COVID-19 period to sharpen inference. We find that results are mainly similar to those presented here, and that our conclusions from the paper remain unchanged.

4.1 Structural Parameter Estimates

Figure 2 displays the previously described prior distributions for the four structural parameters of $\boldsymbol{A}$ in red and the estimated posterior distributions in gray.

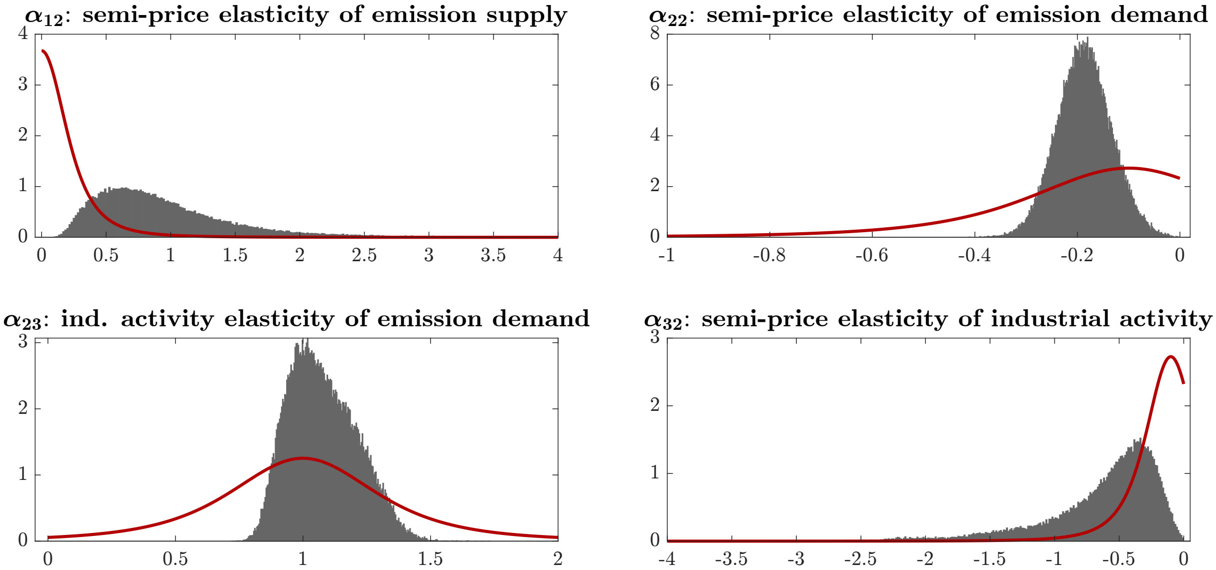

Figure 2: Prior and posterior distributions.

Note: Prior (red) and posterior (gray) distributions of the elements of $\boldsymbol{A}$.

In all four cases, the posterior distributions clearly differ from the prior distributions, indicating that the data is informative about these values. Despite centering the prior for the semi-price elasticity of emission supply α_{12} at zero, the posterior distribution has a location of around 0.6, suggesting that the supply curve is inelastic, but not perfectly inelastic as often assumed. The posterior of the semi-price elasticity of emission demand α_{22} is tightly concentrated around the location of -0.2, which is lower than the prior location of -0.1. The industrial activity elasticity of emission demand α_{23} is much more centered around the location of 1.0 than presumed in the prior. The decreased variance suggests that the data increases the confidence around this parameter value relative to the prior. Finally, the semi-price elasticity of industrial activity is, with a location of about -0.3, lower than the prior location of -0.1.

4.2 The Dynamics of Emission Supply and Demand

Figure 3 shows the cumulated impulse responses of emissions, the EUA price and industrial production (sorted by row) during the 18 months after an emission supply shock, a

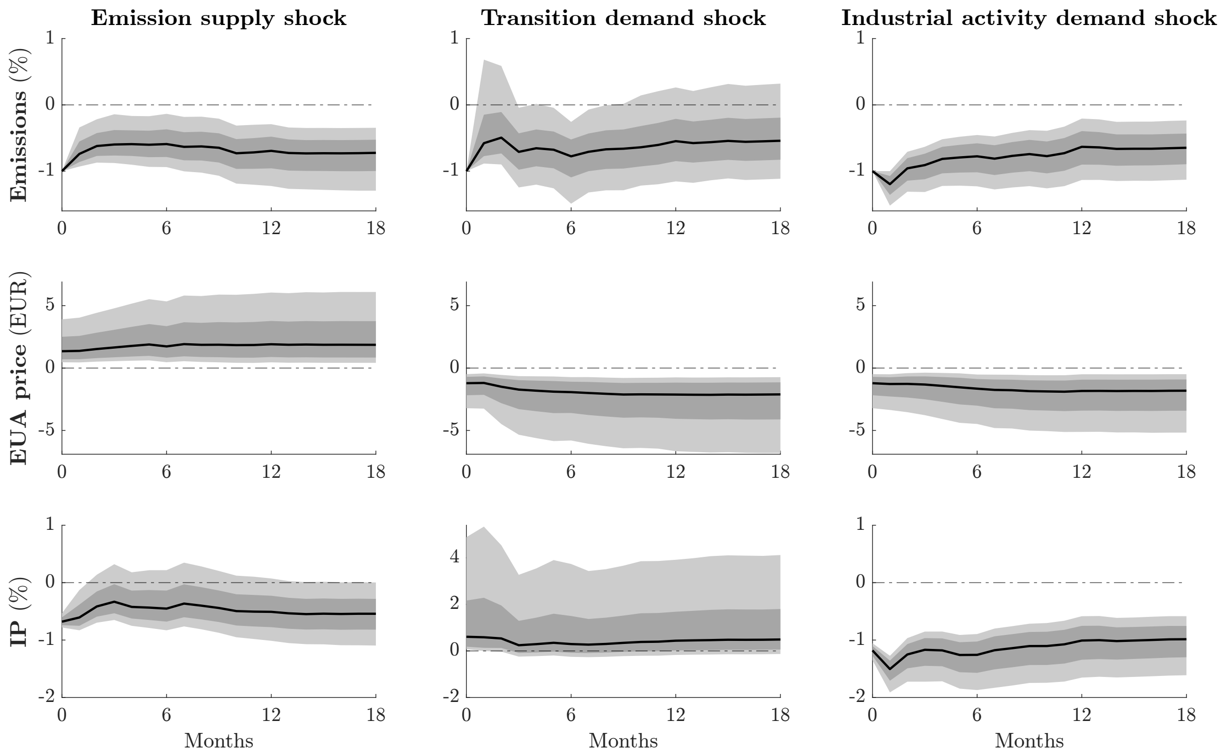

Figure 3: Cumulated impulse responses.

Note: Median cumulated impulse responses are shown in black with the 68% and 90% credible intervals in dark and light gray.

transition demand shock and an industrial activity demand shock (sorted by column). As policymakers are primarily interested in the causes and consequences of emission reductions, we re-scale all three shocks to cause a 1% decrease in emissions on impact. The response is given by the posterior median, while the dark and light gray areas denote the 68% and 90% credible intervals.

The first column shows that an unexpected tightening of emission supply by 1% increases the EUA price by 1.35€ on impact and 1.86€ after one and a half years. This increase in the carbon price leads to higher input costs for firms and dampens industrial production by 0.68% in the short run and 0.54% in the long run. The significant effect on industrial production is different from the evidence in Metcalf and Stock (2023) who find an insignificant (or even slightly positive) effect on carbon taxes on economic activity while it is in line with the results of Känzig (2023). Interestingly, we obtain this result despite allowing α_{23} to be zero.

The second column suggests a transition demand shock that reduces emissions by 1%, elicits a 1.21€ decrease in the EUA price on impact and 2.11€ in the long term. This effect could result from, for example, increased energy generation from renewables, technological progress or a change in consumers and firm behavior. For instance, an unexpectedly large amount of sun, wind or rain or technological progress in harvesting their power would increase energy production from renewable sources, which can be employed at almost zero marginal cost. This results in a substitution away from fossil fuels since they have marginal costs larger than zero, thereby reducing emissions and decreasing the carbon price. At the same time, lower energy and carbon prices significantly increase industrial economic activity in the short run. Similarly, a sudden shift in preferences or unexpected subsidies accelerating the green transition reduces the demand for emissions and the EUA price. Consumers and producers paying voluntarily a premium for low-carbon goods and services and subsidies for low-carbon investments are also positive for producers and will increase industrial production, as depicted in the bottom graph of the second column.

The third column shows that an industrial activity demand shock, which reduces industrial production by 1.18% and emissions by 1% on impact, decreases the EUA price by 1.21€. This result is in line with the observation that emissions and carbon prices under the EU ETS decreased during times of economic crises, such as the GFC, or shutdowns of the industry, as in the recent Covid-19 crisis.

4.3 The Drivers of Carbon Emissions

We now turn to our primary research question and identify the proportion of emission reductions in the European carbon market driven by the supply and two distinct demand factors. Figure 4 shows the median historical decomposition of the cumulated change in emissions compared to December 2004, the month before the start of the EU ETS. The black line shows the percentage change in emissions over this period, whereas the bars depict the relative contribution of each factor in driving emissions. Below we first analyze the role of supply, before turning to the two demand factors.

In essence, the entire emission supply is directly or indirectly regulated by the EU ETS framework, which allows interpreting the contribution of emission supply as the contribution of the EU ETS. The cap and trade system reduced emissions strongly right after its introduction. However, the oversupply of emission allowances towards the end of the first phase nearly canceled the initial effect of the ETS. As discussed in Section 2.2 supply restrictions during 2007 were essentially not binding, and excess allowances could not be transferred to the following period, which is also reflected in the near-zero carbon

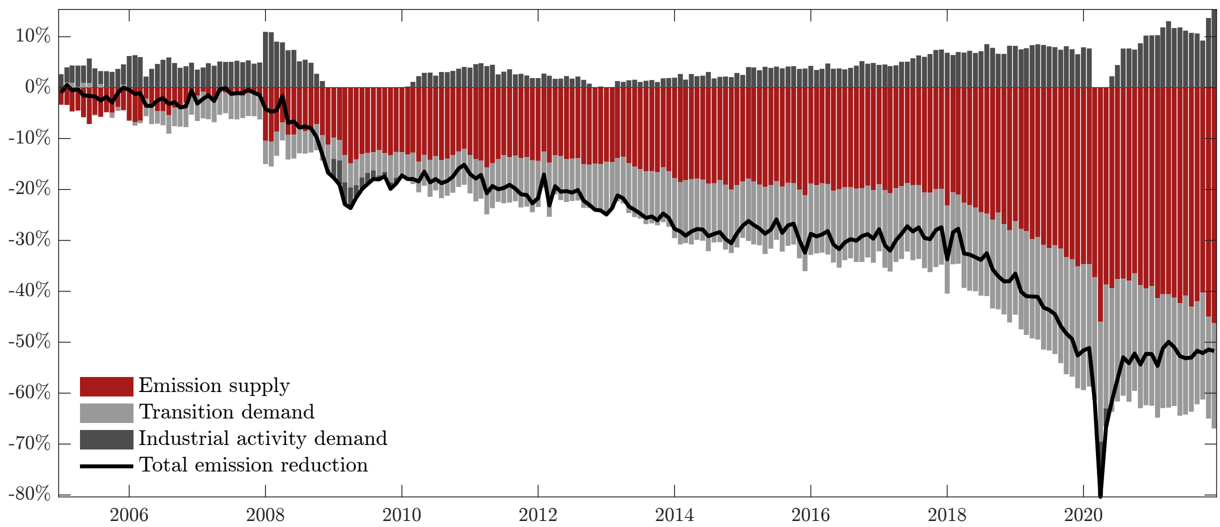

Figure 4: Historical decomposition of the accumulated change in emissions.
Note: The black line shows the percentage change in emissions over this period, whereas the bars depict the relative contribution of each factor in driving emissions.

price at this time. At the beginning of the second phase in 2008, restrictions from the EU ETS were effective again as the EU decreased the cap strongly based on the experience from the pilot phase. Between 2008 and 2017, the EU ETS steadily decreased emissions, totaling to a reduction of 18% in this period. In the following years, the cap and trade system reduced emissions even faster. Until the end of 2021, the EU ETS alone reduced emissions by 46% (in terms of log-growth). Focusing on each phase of the EU ETS, we find, on average, supply shocks contributed to emissions reductions of around 2% in phase one, 14% in phase two, 25% in phase three, and 7% so far in phase four. While the results for phases three and four are new to the literature, our estimates for the first and second phases are close to those reported in a recent study by Dechezleprêtre et al. (2023). The finding that the second phase had a relatively larger impact on emissions reductions compared to the first phase is also in line with other microeconometric studies (e.g., Colmer et al. (2022) and Petrick and Wagner (2014)). Overall, these results suggest that the EU ETS has achieved substantial emission reductions in its covered industrial sectors. Hence, we confirm the presumptions of Stiglitz et al. (2017) that carbon pricing is "an indispensable part of a strategy for reducing emissions". Still, as we turn to now, demand also matters.

The contribution of emission demand to the development of emissions is a two-sided story. First, the contribution of emission demand from industrial economic activity depends strongly on the business cycle. During times of economic boom, industrial activity increases emissions as more fossil fuel-based energy is demanded and more carbon-emitting goods are produced. Conversely, industrial activity reduces emissions during times of crisis due to decreased production and demand for energy, including the GFC, EDC and Covid-19 Crises. The result that economic activity strongly reduced emissions during the GFC is in line with Declercq et al. (2011) and Bel and Joseph (2015). We also find that the magnitude of emission reductions also depends on the severity of the crisis, with industrial activity demand being the primary driver of emissions reductions during the GFC and Covid-19 periods. Overall, demand from industrial activity has increased emissions by 15% since the introduction of the EU ETS.

Second, the historical contribution of transition demand is entirely different as it reduced emissions steadily over time. This observation aligns well with the narrative of capturing factors such as the increased deployment and usage of renewable energy sources and changes in the behavior of consumers, which are expected to evolve relatively slowly and rather steadily over time. Interestingly, the years following the two well-known United Nations Climate Change Conferences, the *Copenhagen Summit* in December 2009 and the *Paris Climate Accords* in December 2015, are characterized by large emission reductions from transition demand. Throughout the considered horizon, transition demand reduced the EU ETS-covered emission by 21%, which is slightly more than the emission demand from industrial economic activity increased emission.

4.4 Transition Demand

To examine the validity of our interpretation of the transition demand factor, we investigate whether its contribution to emission reductions co-varies with the annual share of fossil fuels in the primary energy consumption across the Euro area countries. If our interpretation is valid, the fossil fuel share should correlate closely with emission reductions from transition demand over the sample period. This is exactly what we see in Figure 5. Thus, while we cannot entirely exclude the possibility that other factors may drive transition demand to some extent, we conclude that the shock primarily reflects the transition to a low-carbon economy via a reduction of fossil fuels in the energy mix.

Finally, to examine the importance of transmission demand further, we provide an extension where we use a simplified two-equation version of the model where we collapse transition demand and industrial activity demand into a single emission demand factor.

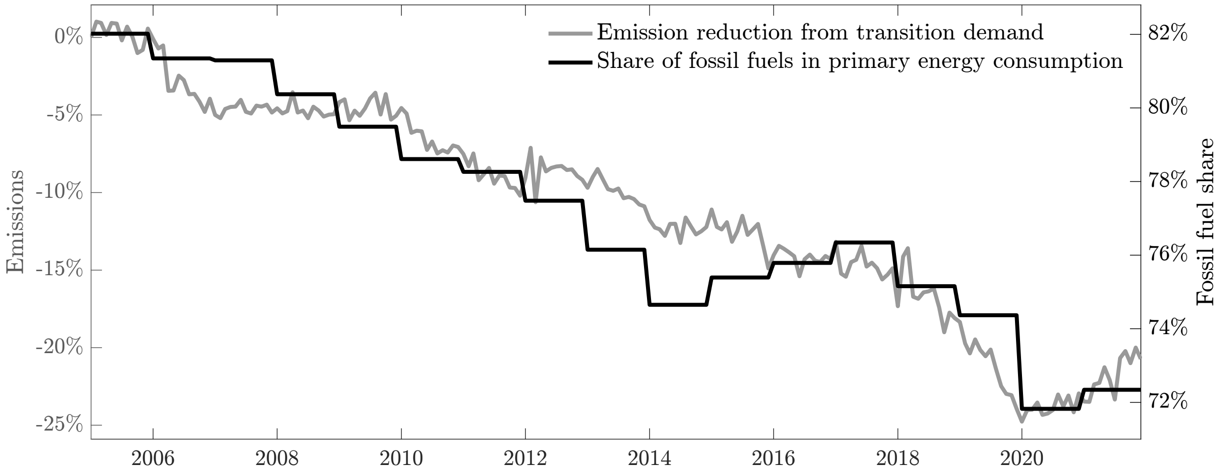

Figure 5: Corroborative evidence for transition demand interpretation.
Note: Contribution of transition demand to emission reductions (gray) and the share of fossil fuels in primary energy consumption (black). The data on energy consumption is obtained from the BP Statistical Review of World Energy 2022 for the Eurozone countries except Malta, which was not available.

Doing so, we find that emission demand plays a relatively larger role than our main results. This difference likely stems from the missing differentiation between the more price-elastic responses of industrial activity and the less price-elastic responses of transition demand. We therefore conclude by considering our three variable setup as the superior model, implying that all three shocks: emission supply shocks, transition demand shocks, and industrial activity demand shocks, matter. Results are in the Online Appendix.

5 Conclusion

In this paper, we introduce an integrated modeling framework for the European carbon market that disentangles the simultaneous forces of emission supply and demand. We begin by showing that the supply curve in the market for emissions is not perfectly inelastic but upward-sloping. This price-sensitivity of emission supply results from several features of the EU ETS regulation, which include the possibility of saving allowances for use or sale in the future, the use of international credits, borrowing allowances from future allocations, and price stabilization mechanisms. As a result, the price and quantity of emissions can be shifted by both emission demand- and supply-side forces. After establishing this fact, we provide a new structural vector autoregressive framework for modeling the simultaneous supply and demand forces in the emissions market. We find that emission supply restrictions set out by the EU ETS have been responsible for the vast majority of emissions reductions. However, we also show that two distinct emission demand factors have played a crucial role. Demand from industrial economic activity has generally increased emissions since the introduction of the EU ETS. However, this increase has been outweighed by other demand factors that reflect primarily the transition to a low-carbon economy.

References

Aastveit, K. A., Bjørnland, H. C., and Cross, J. L. (2023). Inflation expectations and the pass-through of oil prices. *Review of Economics and Statistics*, 105(3):733–743.

Abrell, J., Ndoye Faye, A., and Zachmann, G. (2011). Assessing the impact of the EU ETS using firm level data. Technical report, Bruegel working paper.

Anderson, B. and Di Maria, C. (2011). Abatement and Allocation in the Pilot Phase of the EU ETS. *Environmental and Resource Economics*, 48(1):83–103.

Antolín-Díaz, J. and Rubio-Ramírez, J. F. (2018). Narrative sign restrictions for SVARs. *American Economic Review*, 108(10):2802–2829.

Arias, J. E., Rubio-Ramírez, J. F., and Waggoner, D. F. (2018). Inference based on structural vector autoregressions identified with sign and zero restrictions: Theory and applications. *Econometrica*, 86(2):685–720.

Baranek, B., Boffa, F., and Kastl, J. (2021). Revisiting cap-and-trade in presence of publicly owned polluters: The case of italy 2006-2018. Cepr discussion paper no. dp15989.

Baumeister, C. and Hamilton, J. D. (2015). Sign restrictions, structural vector autoregressions, and useful prior information. *Econometrica*, 83(5):1963–1999.

Baumeister, C. and Hamilton, J. D. (2019). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review*, 109(5):1873–1910.

Baumeister, C. and Hamilton, J. D. (2020). Drawing conclusions from structural vector autoregressions identified on the basis of sign restrictions. *Journal of International Money and Finance*, 109:102250.

Baumeister, C. and Hamilton, J. D. (2022a). Advances in using vector autoregressions to estimate structural magnitudes. *Econometric Theory*, pages 1–39.

Baumeister, C. and Hamilton, J. D. (2022b). Structural vector autoregressions with imperfect identifying information. In *AEA Papers and Proceedings*, volume 112, pages 466–70.

Bayer, P. and Aklin, M. (2020). The European Union emissions trading system reduced CO2 emissions despite low prices. *Proceedings of the National Academy of Sciences*, 117(16):8804–8812.

Bel, G. and Joseph, S. (2015). Emission abatement: Untangling the impacts of the EU ETS and the economic crisis. *Energy Economics*, 49:531–539.

Best, R., Burke, P. J., and Jotzo, F. (2020). Carbon pricing efficacy: Cross-country evidence. *Environmental and Resource Economics*, 77(1):69–94.

Chow, G. C. and Lin, A.-L. (1976). Best linear unbiased estimation of missing observations in an economic time series. *Journal of the American Statistical Association*, 71(355):719–721.

Colmer, J., Martin, R., Muûls, M., and Wagner, U. J. (2022). Does pricing carbon mitigate climate change? Firm-level evidence from the European Union emissions trading scheme.

Dechezleprêtre, A., Nachtigall, D., and Venmans, F. (2023). The joint impact of the European Union emissions trading system on carbon emissions and economic performance. *Journal of Environmental Economics and Management*, 118:102758.

Declercq, B., Delarue, E., and D'haeseleer, W. (2011). Impact of the economic recession on the european power sector's CO2 emissions. *Energy Policy*, 39(3):1677–1686.

Ellerman, A. D. and Buchner, B. K. (2008). Over-allocation or abatement? A preliminary analysis of the EU ETS based on the 2005–06 emissions data. *Environmental and Resource Economics*, 41(2):267–287.

Ellerman, A. D. and McGuinness, M. (2008). CO2 abatement in the UK power sector: evidence from the EU ETS trial period.

European Commission (2022a). Development of EU ETS (2005-2020).

European Commission (2022b). Market Stability Reserve.

Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to the calculations of posterior moments. *Bayesian statistics*, 4:641–649.

Gloaguen, O. and Alberola, E. (2013). Les facteurs explicatifs de l'évolution des émissions de CO2 sur les deux phases de l'EU-ETS: une analyse économétrique. *CDC Climat Recherche Document de travail*, (2013-15).

Jaraite-Kažukauske, J. and Di Maria, C. (2016). Did the EU ETS make a difference? an empirical assessment using Lithuanian firm-level data. *The Energy Journal*, 37(1).

Känzig, D. R. (2023). The unequal economic consequences of carbon pricing. Technical report, National Bureau of Economic Research.

Metcalf, G. E. (2021). Carbon taxes in theory and practice. *Annual Review of Resource Economics*, 13:245–265.

Metcalf, G. E. and Stock, J. H. (2023). The macroeconomic impact of Europe's carbon taxes. *American Economic Journal: Macroeconomics*, 15(3):265–286.

Petrick, S. and Wagner, U. J. (2014). The impact of carbon trading on industry: Evidence from German manufacturing firms. *Available at SSRN 2389800.*

Quilis, E. M. (2021). Temporal disaggregation, https://www.mathworks.com/matlabcentral/fileexchange/69800-temporal-disaggregation. *MATLAB Central File Exchange (Accessed on 12 November 2021).*

Rubio-Ramirez, J. F., Waggoner, D. F., and Zha, T. (2010). Structural vector autoregressions: Theory of identification and algorithms for inference. *The Review of Economic Studies*, 77(2):665–696.

Stiglitz, J. E., Stern, N., Duan, M., Edenhofer, O., Giraud, G., Heal, G. M., La Rovere, E. L., Morris, A., Moyer, E., Pangestu, M., et al. (2017). Report of the high-level commission on carbon prices.

Wagner, U. J., Muûls, M., Martin, R., and Colmer, J. (2014). The causal effects of the European Union Emissions Trading Scheme: evidence from French manufacturing plants. In *Fifth World Congress of Environmental and Resources Economists, Instanbul, Turkey.*

World Bank (2022). Carbon pricing dashboard.

—Online Appendix—

Not for publication

A Data Sources and Analysis

A.1 Sources

Variable	Source	Sample
Emissions		
Verified industrial EU ETS emissions (EA19)	European Environment Agency	2005-2021
Total net emission (UNFCCC, EA19)	European Environment Agency	2003-2020
Unused EUA allowances of (EU ETS area)	European Environment Agency	2005-2020
Subsectoral emissions (NACE2, EA19)	Eurostat	2008-2020
Gross value added		
Gross value added (EA19)	Eurostat	2019
EUA prices		
Various EUA spot and future prices	European Environment Agency	M1/2005-M7/2011
EUA future (ICE-ECX CFI TRc1)	Datastream	M4/2005-M12/2021
EEX primary auction spot prices	European Energy Exchange	M1/2012-M12/2021
Industrial production		
Industrial production index (seasonally and calendar-adjusted, 2015=100, EA19)	Eurostat	M1/2003-M12/2021
Subsectoral industrial production indices (NACE2, seasonally and calendar-adjusted, 2015=100, EA19)	Eurostat	M12/2002-M12/2021
Consumer prices		
HICP for all item (2015=100, EA19)	Eurostat	M1/2003-M12/2021
HICP for energy items (2015=100, EA19)	Eurostat	M1/2003-M12/2021
Fossil fuel consumption		
Oil, natural gas and coal and total primary energy consumption (EA19 without Malta)	BP Statistical Review of World Energy 2022	2005-2021

A.2 Temporal Disaggregation of Emissions

In this Appendix we provide details on the temporal disaggregation procedure used to convert the annual greenhouse gas emissions data to a monthly frequency.

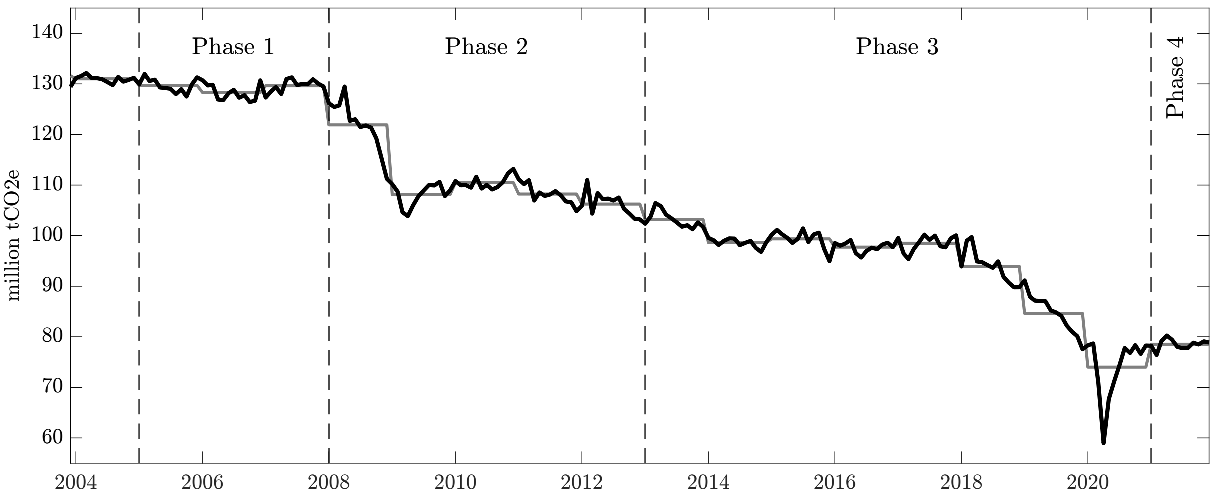

Figure 6: Temporal disaggregation of emissions.
Note: Annual greenhouse gas emissions (gray) and monthly temporally disaggregated greenhouse gas emissions (black).

Temporal Disaggregation. Figure 6 shows the annual series of verified EU ETS emissions of the industry (under the current scope) in gray. To obtain a monthly series of emissions, we follow Känzig (2023) in temporally disaggregating the annual series with the Chow-Lin method provided by the code package of Quilis (2021). By construction, the low-frequency variation of emissions is determined by the annual series. In contrast, the latent high-frequency variation of the temporally disaggregated series is estimated with two monthly indicators that are highly relevant for the development of emissions - the harmonised index of consumer prices for energy items and the industrial production index. While Känzig (2023) uses the usual gross value added-weighted industrial production index, we suggest using an emission-weighted industrial production index for EU ETS-covered sectors to add more precision. The construction of this index is described below.

Emission-weighted industrial production index. The emission-weighted industrial production index is constructed using six subsectors that are covered by the EU ETS. Together, these six subsectors account for 89.2% of the industry's emissions but only for 24.5% of the industry's gross value added in 2019. The emission weights for the index construction are based on the average share of emissions of the subsectors across the years 2008 to 2020, the years for which data on subsectoral emissions was available.

Subsector	NACE Code	Share of industry emissions in 2019	Share of industry gross value added in 2019	Index emission weight
Manuf. of paper and paper products	C17	2.2%	1.9%	2.3%
Manuf. of coke and refined petroleum products	C19	8.8%	0.9%	8.6%
Manuf. of chemicals and chemical products	C20	9.6%	5.9%	9.6%
Manuf. of other non-metallic mineral products	C23	11.8%	2.9%	12.2%
Manuf. of basic metals	C24	10.6%	2.7%	10.7%
Electr., gas, steam and air cond. supply	D or D35	46.2%	10.2%	56.5%
Total	-	89.2%	24.5%	100%

Table 1: Subsectors and index weights.

Note: Considered industrial subsectors used to construct emission-weighted industrial production index for the EU ETS.

Figure 7 provides a comparison of the original gross value added-weighted industrial production index and our emission-weighted industrial production index for the EU ETS-covered sectors.

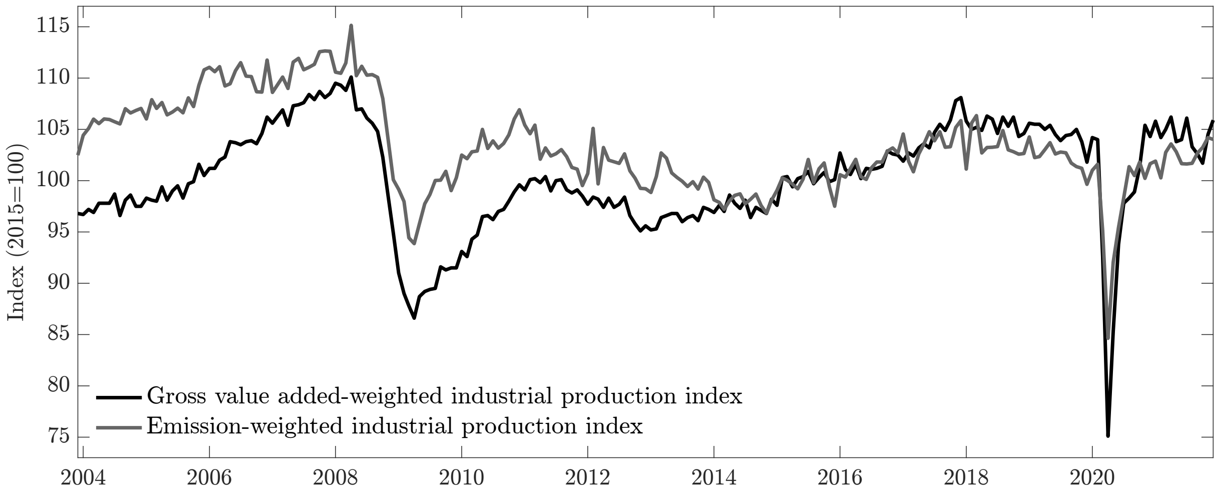

Figure 7: Comparison of the two industrial production indices.

Note: While Känzig (2023) uses the original gross value added-weighted industrial production index, we use our constructed emission-weighted industrial production index to temporally disaggregate the annual verified industrial emissions.

A.3 EUA Spot and Future Prices

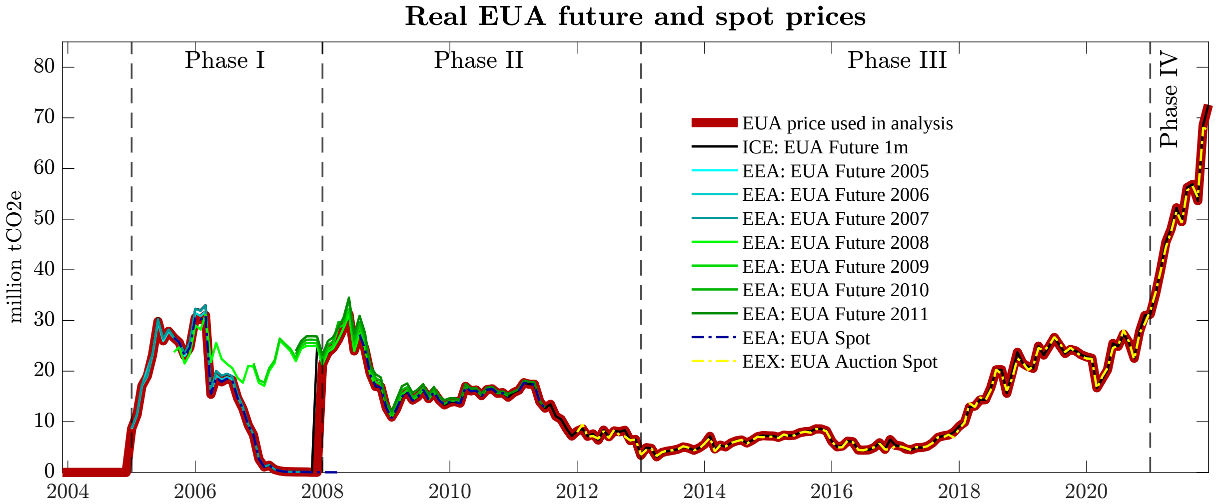

Figure 8: Constructed EUA price in comparison to available spot and future prices.

Figure 8 displays our constructed EUA price (thick red line) in comparison to EUA (auction) spot and future prices. Ideally, we would use end-of-month spot prices to represent the current price of EUAs in each month. Unfortunately, spot prices (dashed-dotted lines) are not available for the whole period. Moreover, spot prices from auctions (yellow dashed-dotted line, obtained from European Energy Exchange) are not auctioned every day. Hence, some end-of-month spot prices are rather mid-month prices. For example, auctions in December end roughly one week before Christmas. Hence, we rely on the Intercontinental Exchange's (ICE) Futures 1-month ahead from the second phase on as they are frequently traded and hence, they provide reliable end-of-month EUA prices.

In general, EUA (auction) spot and 1-month ahead or current-year future prices co-vary closely. Future prices several years ahead are typically slightly above future prices for the current years, which represents the markup for future price certainty. Moreover, EUA futures for the second phase of the EU ETS (green lines, Futures 2008-2011) differ from EUA prices of the first phase (turquoise lines, Futures 2005-2007) as EUAs of the first phase could not be transferred to the second phase.

B Additional Results

B.1 Prior Distributions

Parameter	Description	Restriction	Location	Scale	df
α_{12}	semi-price elasticity of emission supply	≥ 0	0	0.2	3
α_{22}	semi-price elasticity of emission demand	≤ 0	-0.1	0.2	3
α_{23}	industrial activity elasticity of emission demand	≥ 0	1.0	0.3	3
α_{32}	semi-price elasticity of industrial activity	≤ 0	-0.1	0.2	3

Table 2: Prior distributions on the elements of matrix $\boldsymbol{A}$.

Note: In each case we specify Student's $t(\mu, \sigma, \nu)$ distributions in which μ is the location parameter, σ is the scale parameter, and ν is the degree of freedom parameter.

B.2 Convergence Diagnostics

We present evidence for the convergence of the parameter draws with the help of two diagnostic tools. Figure 9 shows the trace for the sample draws of the four structural parameters α_{12}, α_{22}, α_{23} and α_{32} of matrix $\boldsymbol{A}$, suggesting that the algorithm sufficiently explores the parameter space with a high degree of statistical efficiency.

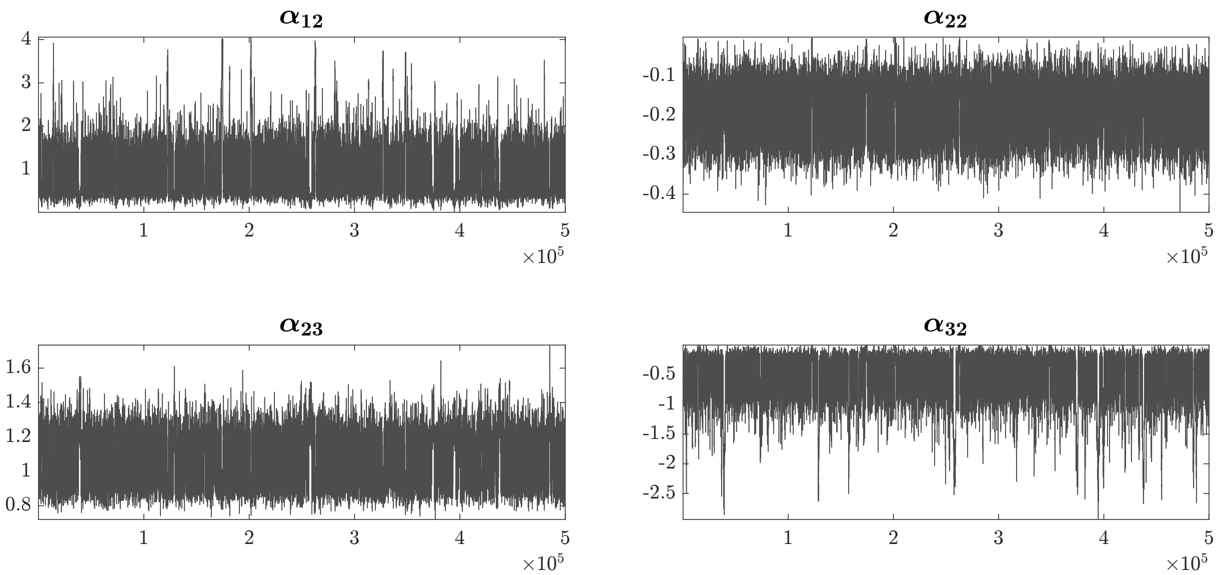

Figure 9: Trace of the sample draws.
Note: Shows the sample draws of the structural parameters of matrix $\boldsymbol{A}$ after discarding the burn-in.

Table 3 shows the results for the convergence diagnostic test of Geweke (1992) using the first 10% and last 50% of the sample parameter draws that are used for inference. The results suggest that the posterior means of all four parameters are stable across the draws.

	χ^2 probability for equality of means			
Decay in autocov. function	α_{12}	α_{22}	α_{23}	α_{32}
4% taper	0.615	0.908	0.441	0.638
8% taper	0.624	0.912	0.473	0.663
15% taper	0.611	0.911	0.461	0.650

Table 3: Convergence diagnostics following Geweke (1992).

B.3 Historical Decomposition of the EUA Price Development

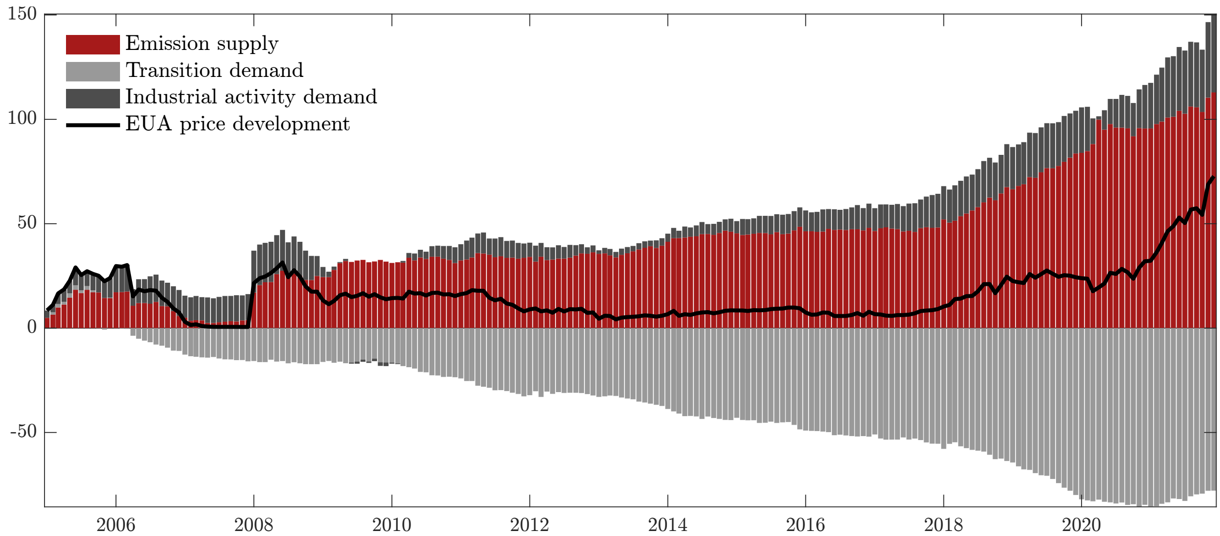

Figure 10: Historical decomposition of the accumulated change in the EUA price. Note: Accumulated change is denoted in Euro per EUA.

Figure 10 shows the historical contribution of emission supply (red), transition demand (light gray) and industrial economic activity (dark gray) to the development of the EUA price (black line). It shows that the EU ETS limiting emission supply is by far the largest contributor, but emission demand also plays a substantial role. Emission demand due to increased industrial economic activity has driven up the EUA price. For example, it built up pressure on the EUA price during the boom before the financial crisis, but reduced pressure on the EUA price during the financial crisis and during the following Euro crisis. In contrast, transition demand has decreased the EUA price slowly and continuously over time. The latter finding is also in line with the interpretation of a rather slow and steady deployment of renewables and slow changes in consumer behavior.

B.4 Shock Series

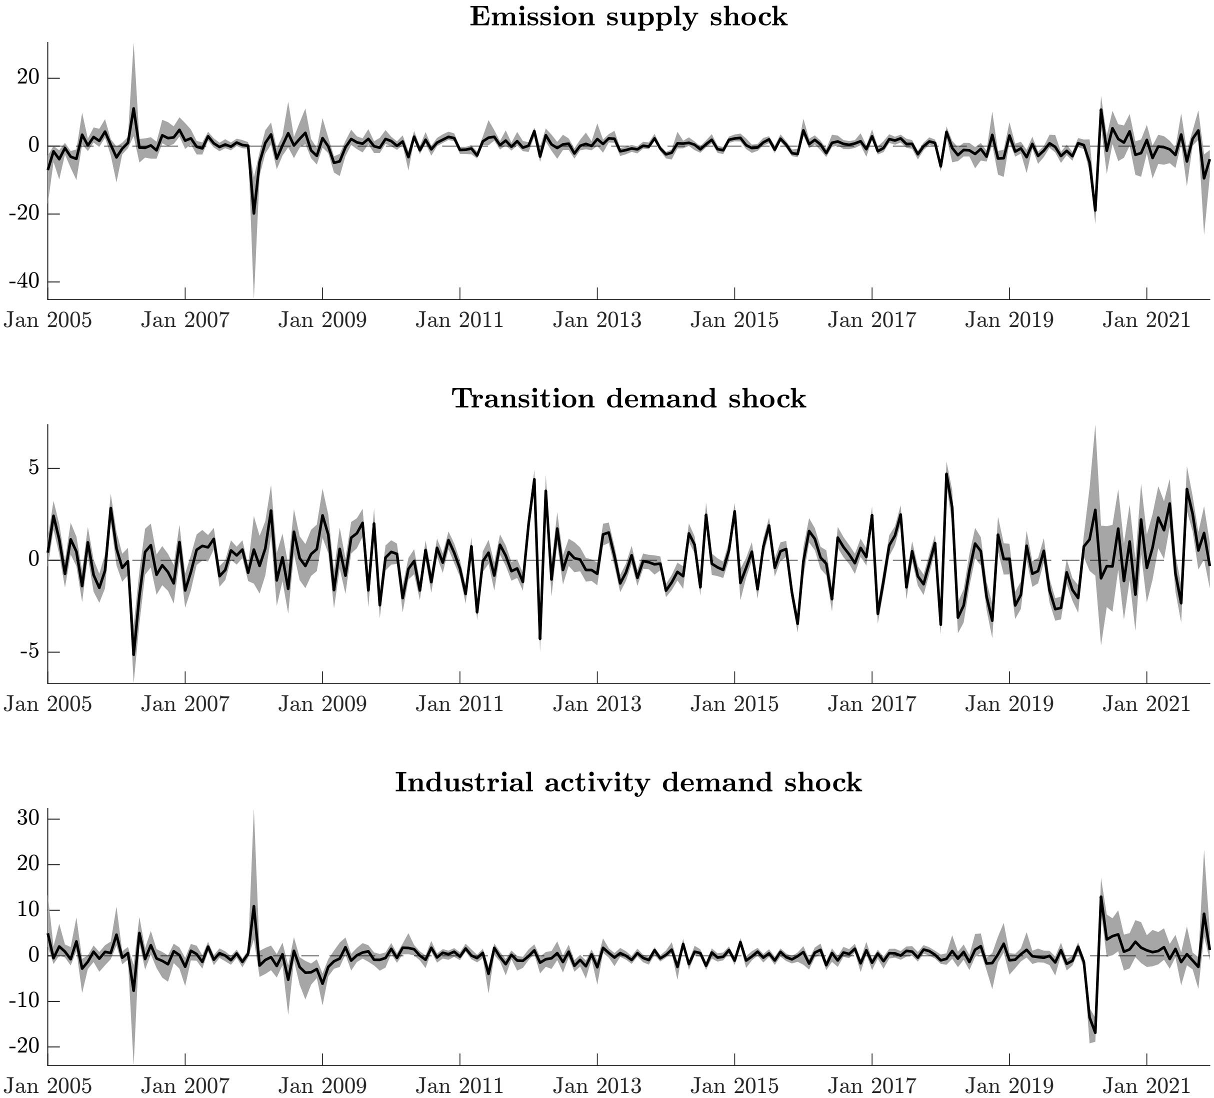

Figure 11: Shock series.

Note: Each graph reports the median shock series (black line) and their credible intervals (gray area).

C Robustness checks

C.1 Robustness to Lag Length

We check for robustness with regard to the number of lags included in the equations (1)-(3). Instead of twelve lags, we control for six or eighteen lags. All other settings remain as in the main analysis. Overall, we find that the lag length influences neither our results nor their interpretation.

C.1.1 Six Lags

Figure 12 shows the posterior distributions of the structural parameters, which are very close in location and shape to the posterior distributions of the main results.

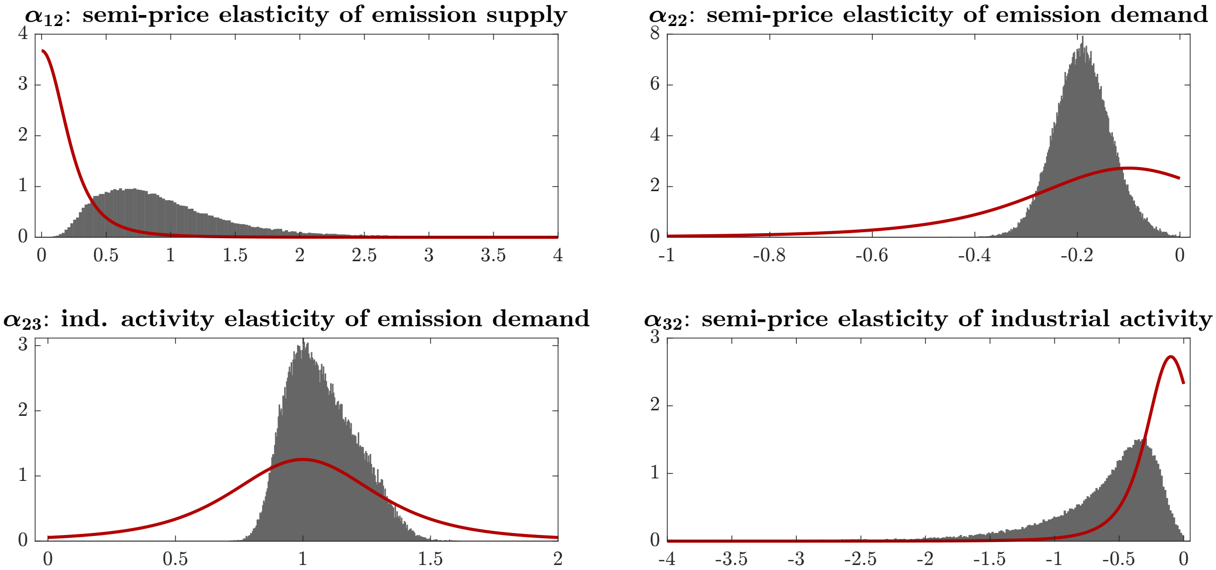

Figure 12: Robustness test with six lags: prior and posterior distributions.

Figure 13 shows the trace of the draws of the structural parameters after discarding the burn-in, indicating that the MH MCMC sufficiently explores the parameter space with a high degree of statistical efficiency.

Figure 14 displays the median impulse responses and the credible intervals, which are very similar in shape and magnitude to those of the main results.

Figure 15 shows the historical decomposition of emissions, which is similar to main results.

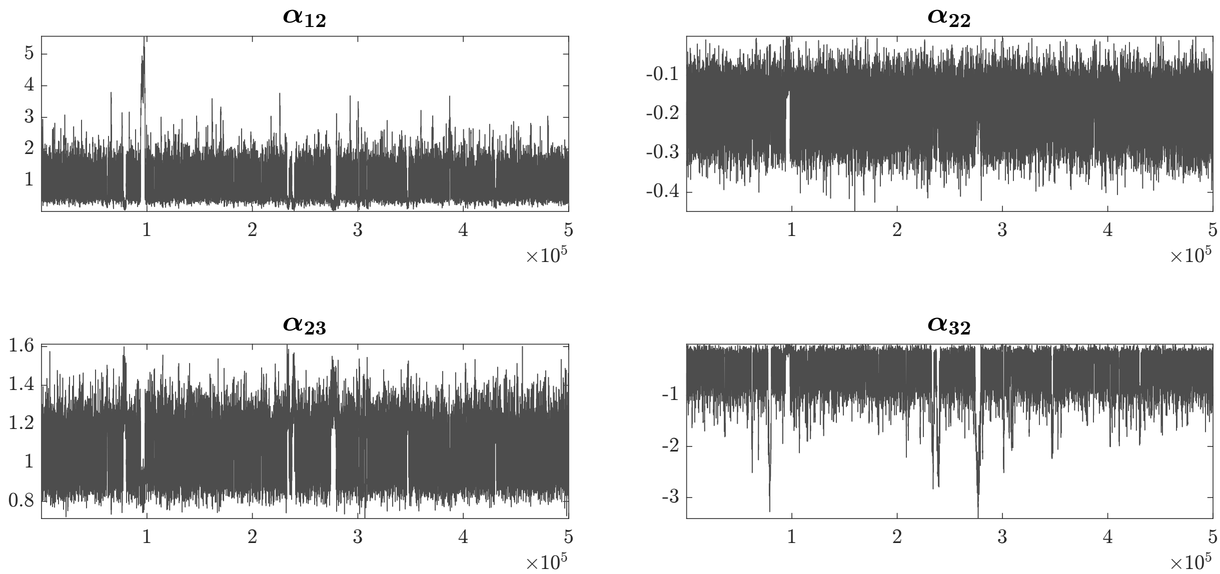

Figure 13: Robustness test with six lags: Trace of the sample draws.

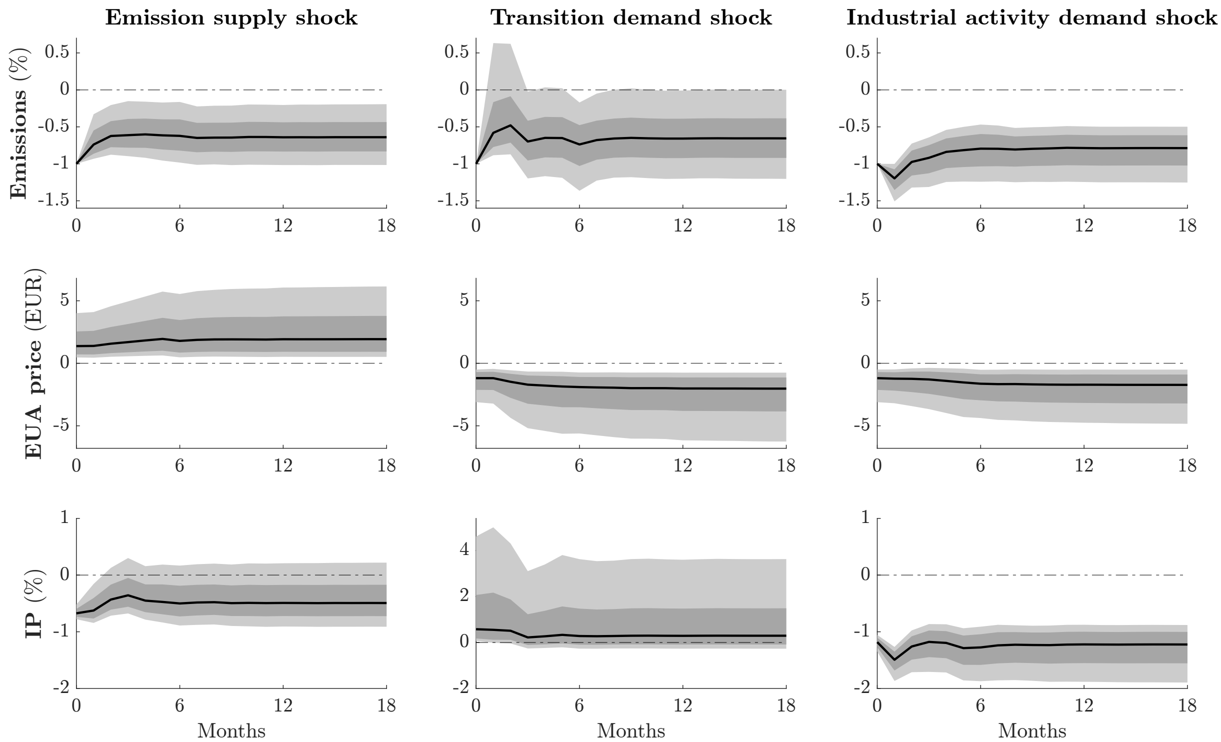

Figure 14: Robustness test with six lags: Accumulated impulse responses.

Note: Graphs display the median accumulated impulse responses (black) and their 68% and 90% credible intervals (dark and light gray).

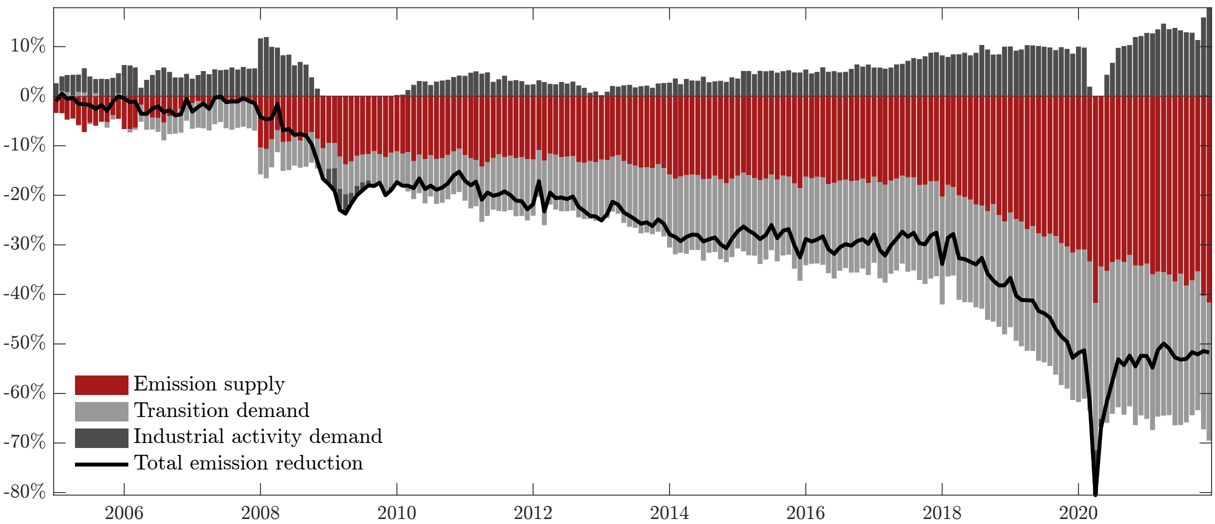

Figure 15: Robustness test with six lags: Historical decomposition of the accumulated change in emissions.

C.1.2 Eighteen Lags

Figure 16 shows the posterior distributions of the structural parameters, which are very close in location and shape to the posterior distributions of the main results.

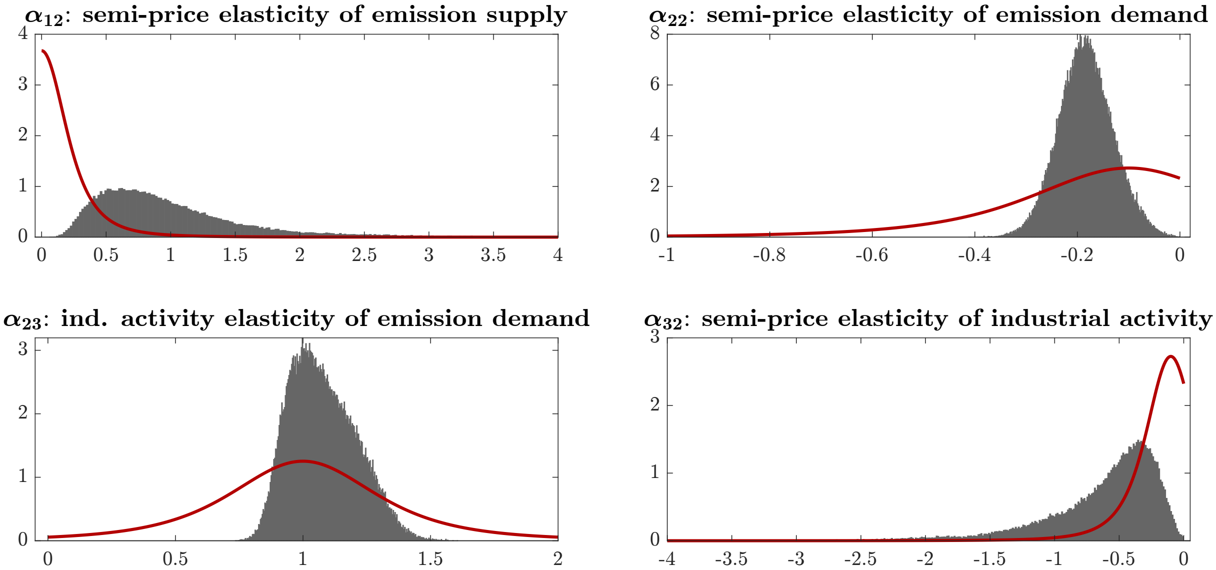

Figure 16: Robustness test with eighteen lags: prior and posterior distributions.
Note: Posterior distributions of the elements of $\boldsymbol{A}$ are shown after discarding the burn-in.

Figure 17 shows the trace of the draws of the structural parameters after discarding the burn-in, suggesting that the algorithm sufficiently explores the parameter space with a high degree of statistical efficiency.

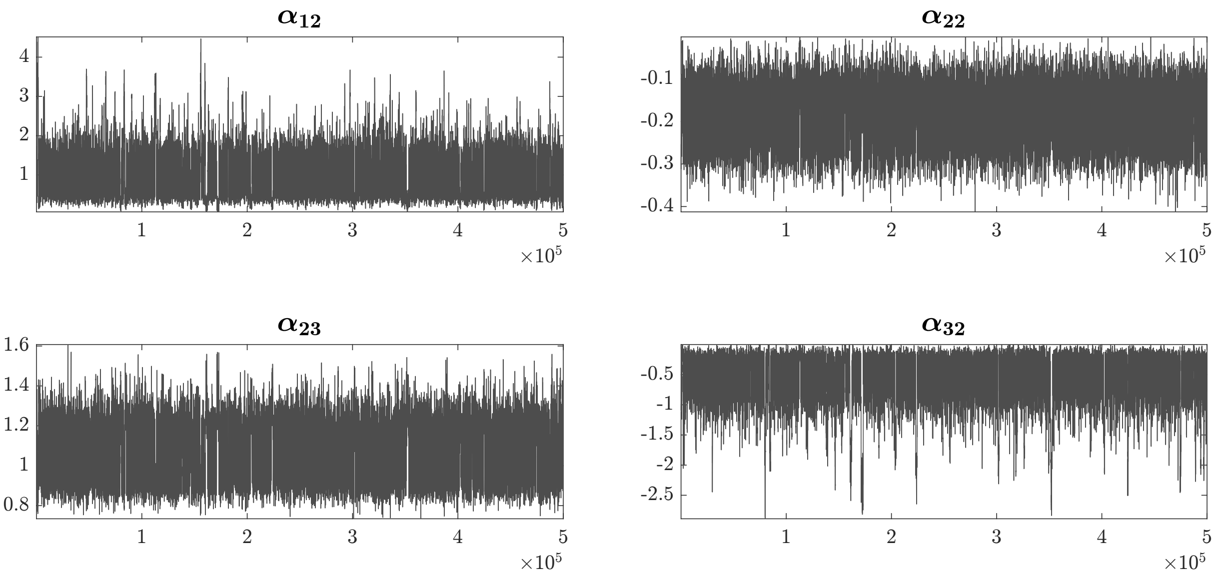

Figure 17: Robustness test with eighteen lags: Trace of the sample draws.
Note: Sample draws for the matrices $\boldsymbol{A}$ after discarding the burn-in.

Figure 18 displays the median impulse responses and the credible intervals, which are very similar in shape and magnitude to the main results.

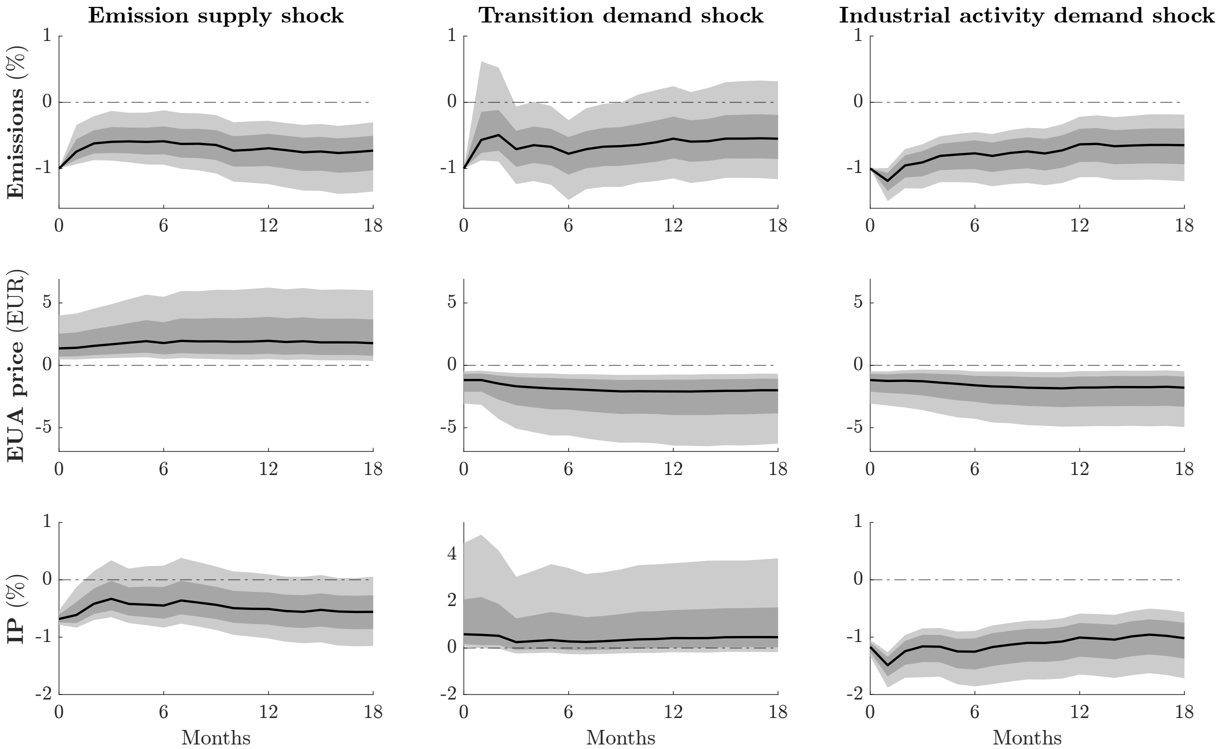

Figure 18: Robustness test with eighteen lags: Accumulated impulse responses. Note: Graphs display the median accumulated impulse responses (black) and their 68% and 90% credible intervals (dark and light gray).

Figure 19 shows the historical decomposition of emissions, which is very similar to main results.

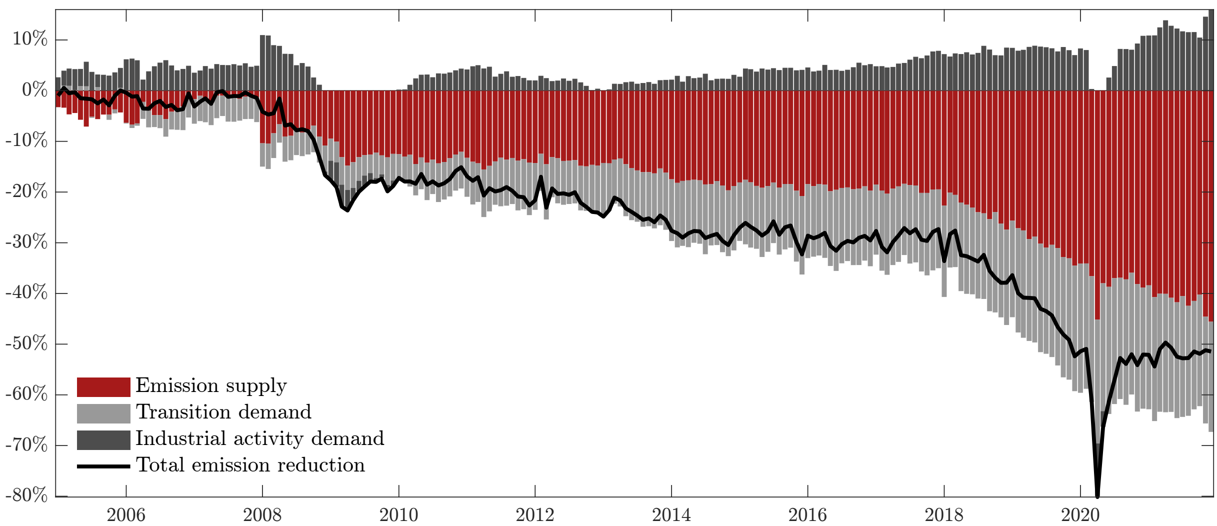

Figure 19: Robustness test with eighteen lags: Historical decomposition of the accumulated change in emissions.

C.2 Robustness to the Covid-19 Period

To check for robustness regarding the recent Covid-19 crisis, we end the sample in December 2019 and repeat the empirical analysis with otherwise same settings. All of our drawn conclusions again remain.

Figure 20 shows the posterior distributions of the structural parameters, which are similar to those found with the full sample.

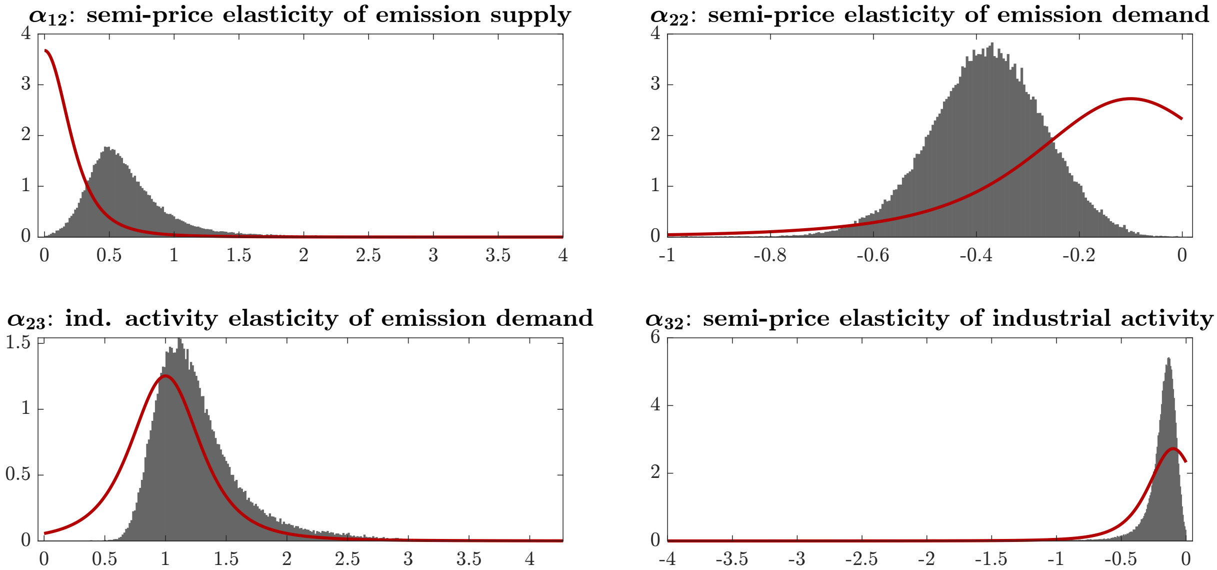

Figure 20: Robustness test excluding Covid-19 period: prior and posterior distributions.
Note: Posterior distributions of the elements of $\boldsymbol{A}$ are shown after discarding the burn-in.

Figure 21 shows the trace of the draws of the structural parameters after discarding the burn-in, suggesting that the algorithm sufficiently explores the parameter space with a high degree of statistical efficiency.

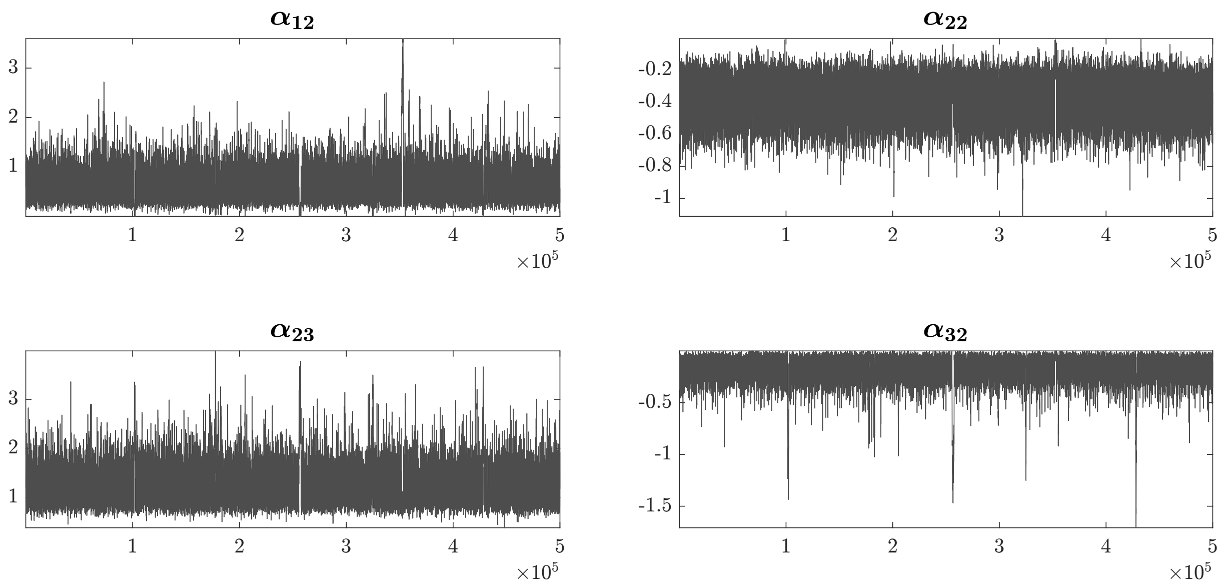

Figure 21: Robustness test excluding Covid-19 period: Trace of the sample draws.
Note: Sample draws for the matrices $\boldsymbol{A}$ after discarding the burn-in.

Figure 22 displays the median impulse responses and the credible intervals. While some minor differences in magnitude can be observed, the overall inference is in line with our main results.

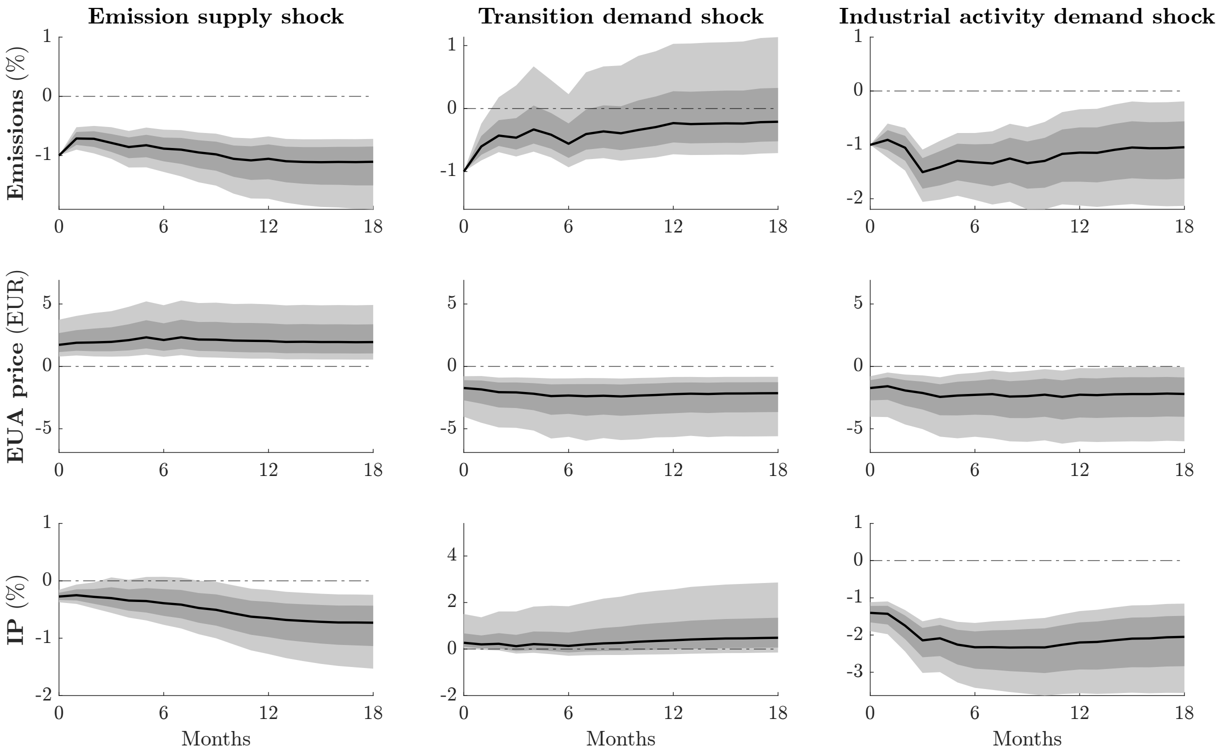

Figure 22: Robustness test excluding Covid-19 period: Accumulated impulse responses.
Note: Graphs display the median accumulated impulse responses (black) and their 68% and 90% credible intervals (dark and light gray).

Figure 23 shows the historical decomposition of emissions, suggesting the contribution of the EU ETS was more significant in the pre-Covid-19 sample than in the full sample.

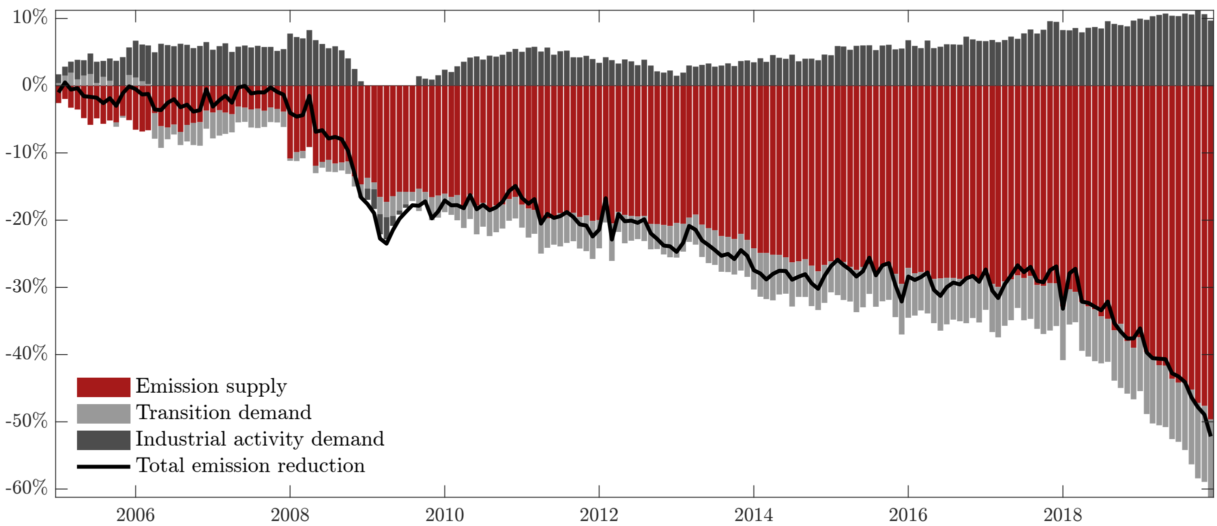

Figure 23: Robustness test excluding the Covid-19 period: Historical decomposition of the accumulated change in emissions.

C.3 A Two-Variable System Without Industrial Production

To provide an impression of how the results would look like if we collapsed industrial activity demand and transition demand to only one demand factor, *emission demand*, we leave out industrial production from the model by removing equation (3) from the system of equations and removing y_t from equation (2). This changes the interpretation of the shock in equation (2), which now captures unexpected changes of all emission demand factors. The two variable setup therefore provides an impression of the additional insights that we gain from splitting general emission demand into two main emission demand factors.

Figure 24 shows the posterior distributions of the structural parameters. While the posterior of α_{12} is similar in location to the main results of the three-variable system but has a slightly larger shape, α_{22} changes considerably and is rather similar to α_{32} of the three-variable system. This is little surprising as we expect that most of the price elasticity in emission demand is due to price adjustments in industrial economic activity.

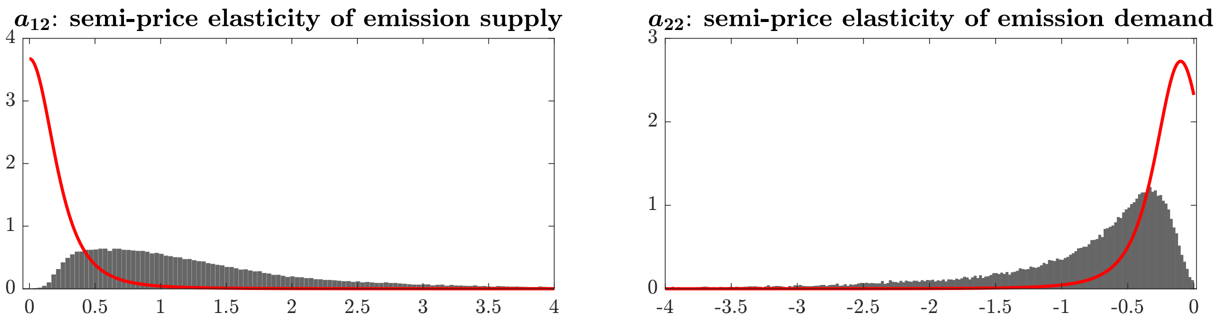

Figure 24: Robustness test for two-variable system: prior and posterior distributions.
Note: Posterior distributions of the elements of $\boldsymbol{A}$ are shown after discarding the burn-in.

Figure 25 shows the trace of the draws of the structural parameters after discarding the burn-in, suggesting that the algorithm sufficiently explores the parameter space with a high degree of statistical efficiency.

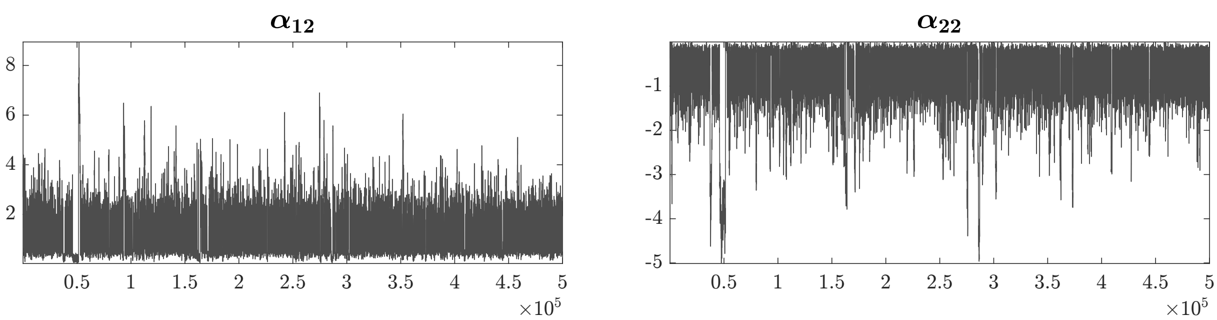

Figure 25: Robustness test for two-variable system: Trace of the sample draws.
Note: Sample draws for the matrices $\boldsymbol{A}$ after discarding the burn-in.

Figure 26 displays the median impulse responses and the credible intervals. The impulse responses of the emission supply shock are similar in magnitude to those of the three-variable system, but uncertainty around the median response increased considerably. The impulse responses of the emission demand shock are mostly similar to those of the industrial activity demand shock from the three-variable system, which confirms the idea that industrial activity shocks are the most important emission demand shocks.

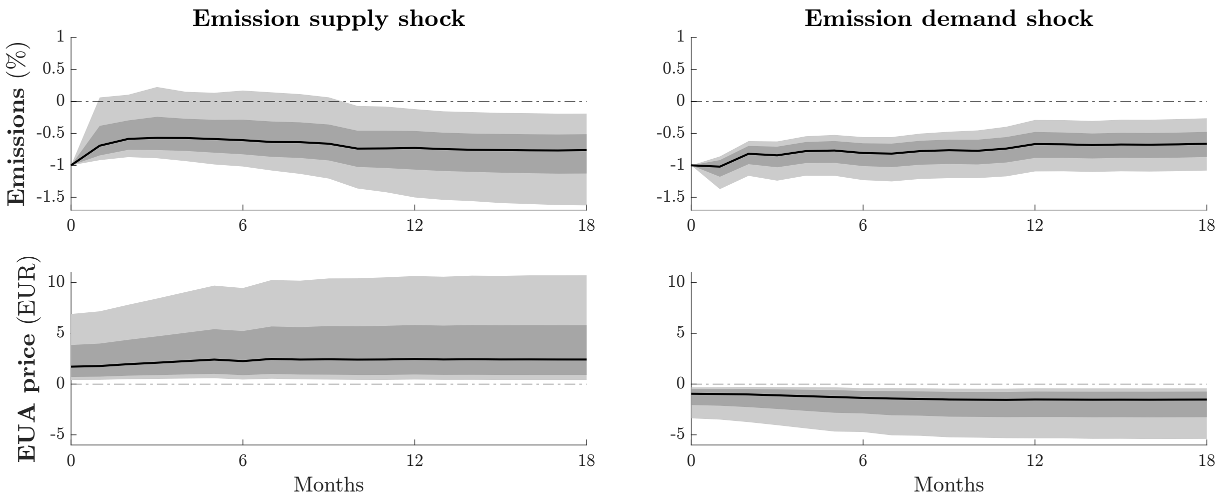

Figure 26: Robustness test for two-variable system: Accumulated impulse responses.
Note: Graphs display the median accumulated impulse responses (black) and their 68% and 90% credible intervals (dark and light gray).

Figure 27 shows the historical decomposition of the development of emissions, which suggests that emission demand plays a more important role than in the three-variable system. Moreover, note that emission demand in the two-variable system reflects the variation of transition demand plus industrial activity demand. Thus, this difference likely stems from the missing differentiation between the more price-elastic responses of industrial activity and the less price-elastic responses of transition demand. We therefore consider our three variable setup as the more accurate and informative choice.

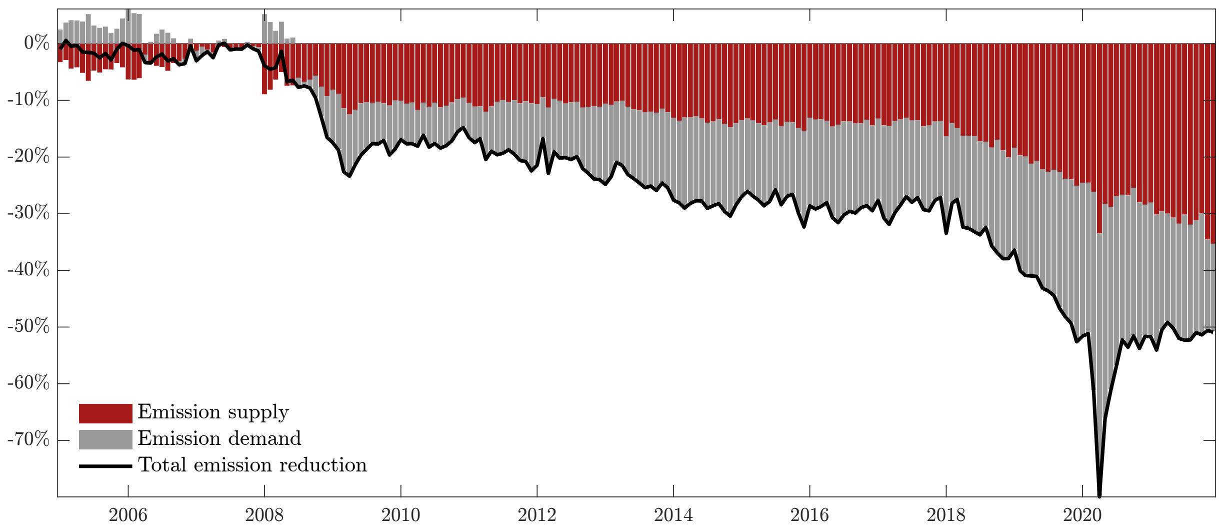

Figure 27: Robustness test for two-variable system: Historical decomposition of the accumulated change in emissions.

C.4 Narrative restrictions

To investigate whether additional assumptions can sharpen inference we also consider an additional experiment in which we introduce a set of narrative restrictions during the COVID-19 period, which is defined here to be March to April of 2020 since this covers the large decrease in industrial production (see Figure 1). Specifically, we assume that the industrial activity demand shock was the dominant shock during this period. This is done by restricting the set of admissible structural models to be those in which the industrial activity demand shock accounts for at least half of the total variation across the structural shock series in this period. Given that there were no changes in ETS policy during this period we view these restrictions as rather mild. The restrictions are implemented with a straightforward adaptation of the narrative sign restrictions algorithm in Antolín-Díaz and Rubio-Ramírez (2018) to our framework. The main difference is that we implement the sign restrictions using the Bayesian methods developed in Baumeister and Hamilton (2015) and Baumeister and Hamilton (2019), while Antolín-Díaz and Rubio-Ramírez (2018) follow the Bayesian methods developed in Rubio-Ramirez et al. (2010) and Arias et al. (2018).

The estimated posterior distributions are in Figure 28. Relative to the main results, we find that the posterior distribution of the semi-price elasticity of emission supply (α_{12}) has the same shape, but is slightly shifted to the right. The posterior median of this distribution is around 1.1, suggesting that the supply curve is elastic, as opposed to perfectly inelastic, as is often assumed. The posterior distribution of the semi-price elasticity of emission demand (α_{22}) is unaffected by these restrictions, remaining tightly concentrated around a location of -0.2. The location of the posterior distribution of the elasticity for industrial activity of emission demand (α_{23}) is also unchanged (located at 1.0), however the distribution here is much more tightly located around this value relative to the main results. Finally, the semi-price elasticity of industrial activity is located around -0.5, which is slightly lower than the prior location of -0.3 found in the main results. The main reason for this is the increased mass in the left tail of the distribution under the narrative restrictions.

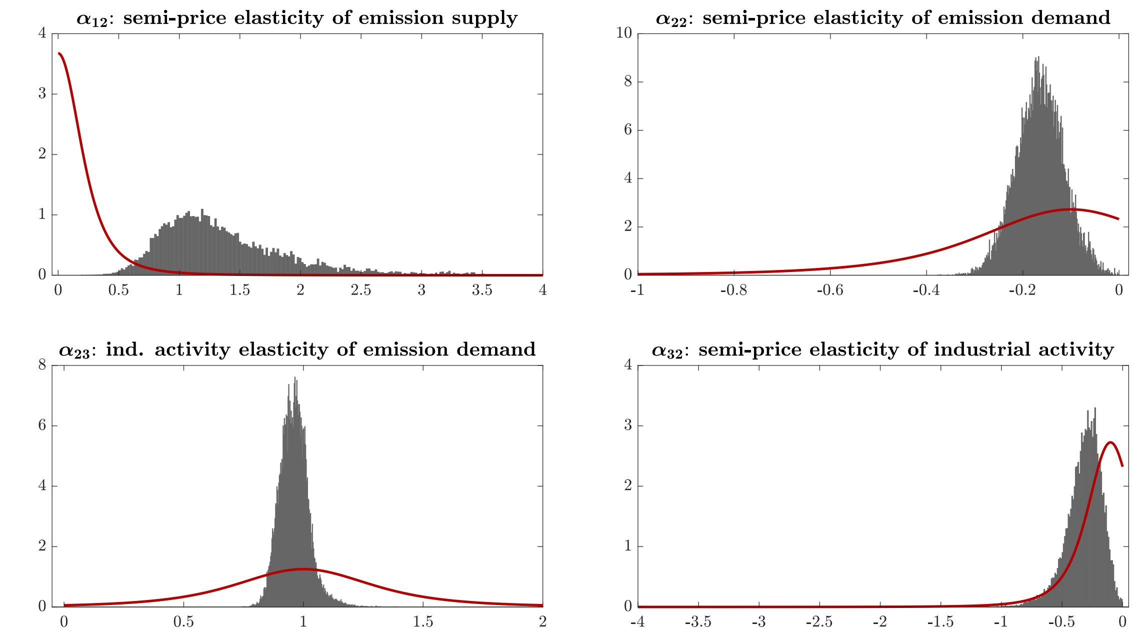

Figure 28: Prior and posterior distributions.

Note: Prior (red) and posterior (gray) distributions of the elements of $\boldsymbol{A}$.

The IRFs are reported in Figure 29. Despite some changes in the posterior distributions of the parameters, the posterior medians of the IRFs are identical to the main results. The only difference between these results and those in the main text is that these IRFs have smaller credible sets, which is indicative of less estimation uncertainty. Thus, all of our main conclusions regarding the IRfs remain unchanged.

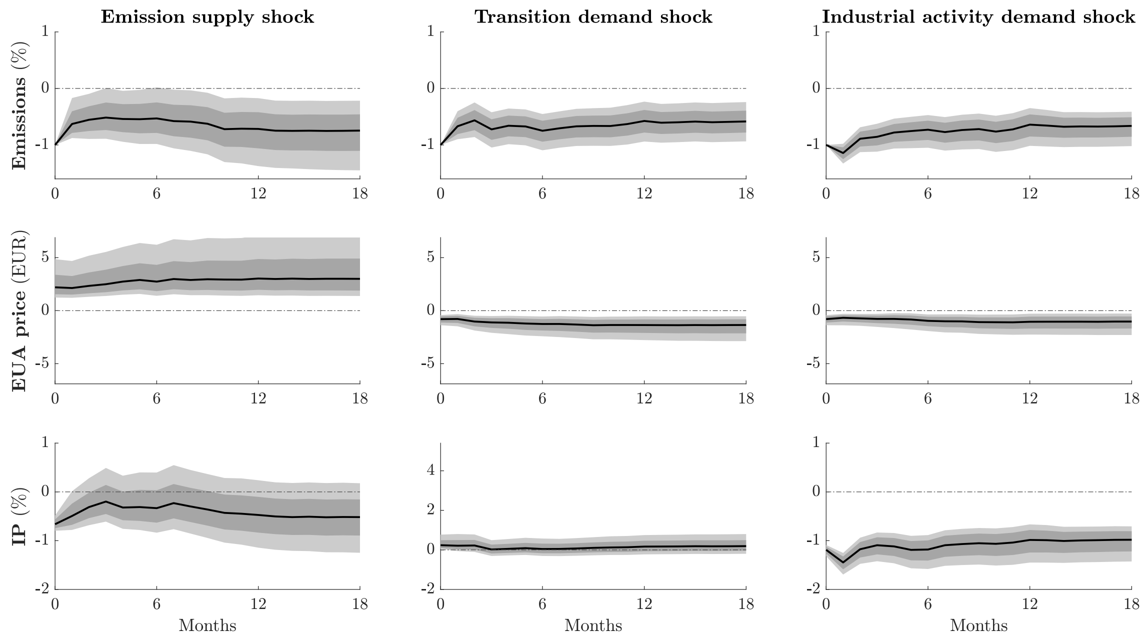

Figure 29: Cumulated impulse responses.

Note: Median cumulated impulse responses are shown in black with the 68% and 90% credible intervals in dark and light gray.

Figure 30 shows the median historical decomposition of the cumulated change in emissions. The general patterns of the contributions are the same as the main results, however the size of the contribution provided by the supply shocks are smaller, relative to the main results, while the contribution of the transition demand is larger.

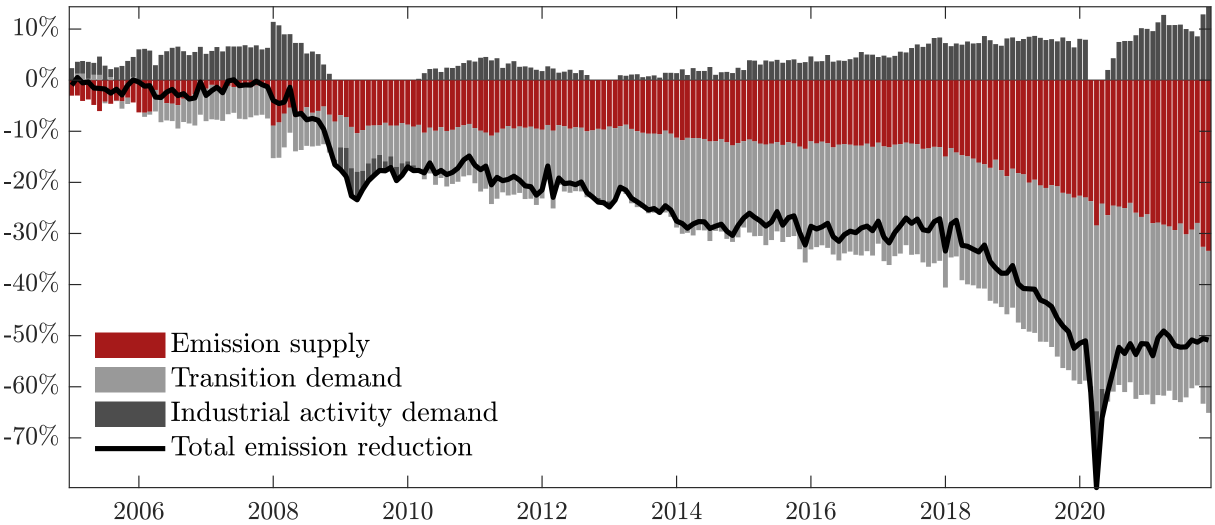

Figure 30: Historical decomposition of the accumulated change in emissions.
Note: The black line shows the percentage change in emissions over this period, whereas the bars depict the relative contribution of each factor in driving emissions.

Figure 31 shows the structural shocks. The main difference compared to the main shock series is that, by assumption, the emissions supply shocks are smaller during the COVID-19 period. One other notable difference is that the transition demand shock is also smaller during this period.

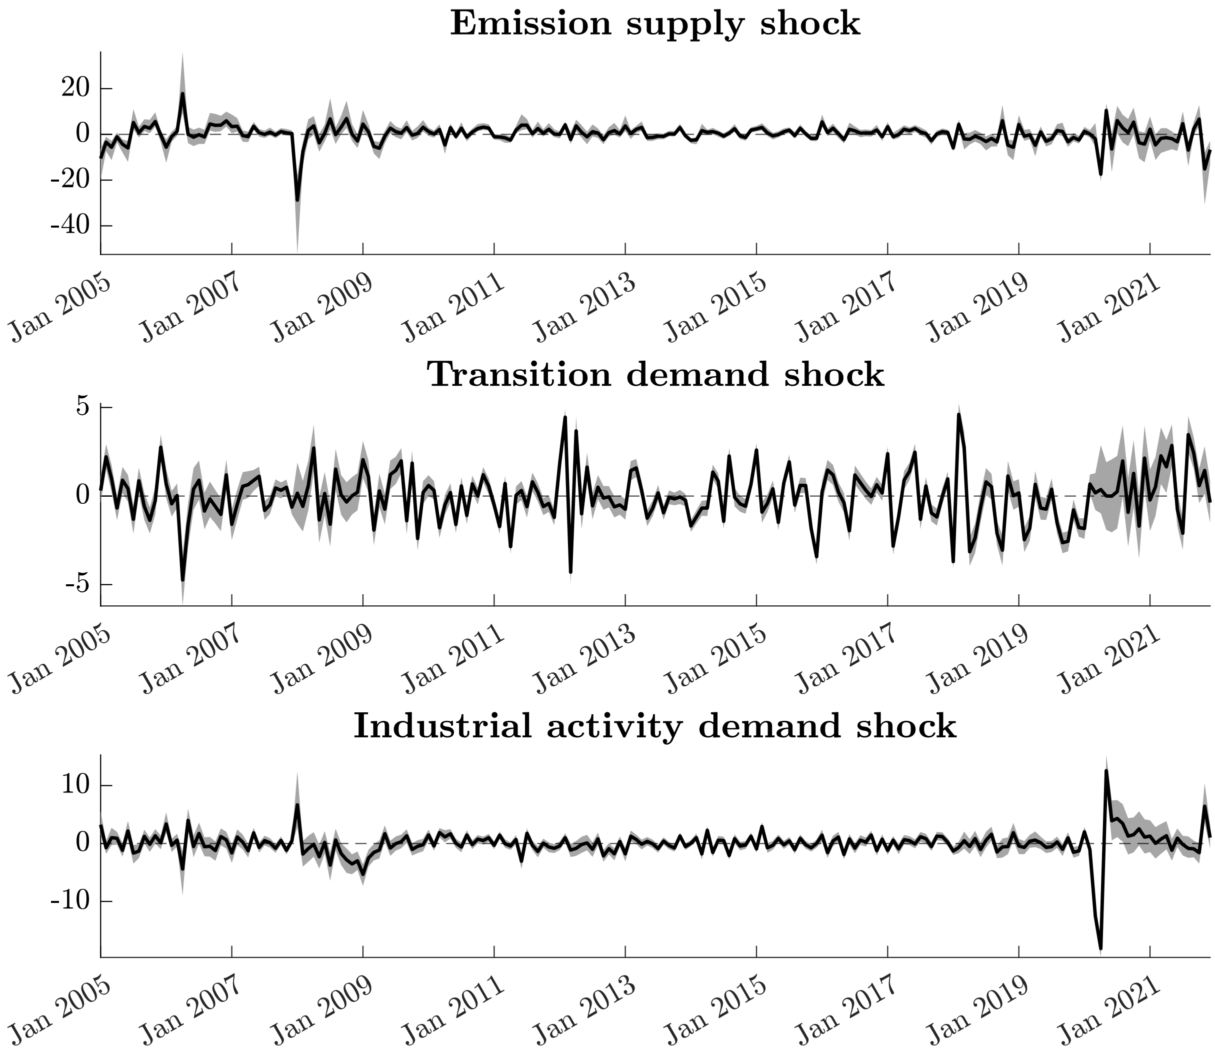

Figure 31: Shock series.

Note: Each graph reports the median shock series (black line) and their credible intervals (gray area).