



**CEEPR**

**Center for Energy and Environmental Policy Research**

**Designing a US Market for CO<sub>2</sub>**

by

**John E. Parsons, A. Denny Ellerman and  
Stephan Feilhauer**

**09-001**

**January 2009**

**A Joint Center of the Department of Economics, Laboratory for Energy  
and the Environment, and Sloan School of Management**



## Designing a US Market for CO<sub>2</sub>

John E. Parsons<sup>\*</sup>, A. Denny Ellerman<sup>\*</sup> and Stephan Feilhauer<sup>\*\*</sup>

January 2009

*In this paper we focus on one component of the cap-and-trade system: the markets that arise for trading allowances after they have been allocated or auctioned. The efficient functioning of the market is key to the success of cap-and-trade as a system. We review the performance of the EU CO<sub>2</sub> market and the US SO<sub>2</sub> market and examine how the flexibility afforded by banking and borrowing and the limitations on banking and borrowing have impacted the evolution of price in both markets. While both markets have generally functioned well, certain episodes illustrate the importance of designing the rules to encourage liquidity in the market.*

The United States may soon have a market for carbon. President-elect Obama has expressed his support for a cap-and-trade system, targeting a reduction in CO<sub>2</sub> emissions to 1990 levels by 2020 and reducing them an additional 80% by 2050. CFTC Commissioner Bart Chilton forecasted this past summer that “Even with conservative assumptions, this could be a \$2 trillion futures market in relatively short order.” Legislation for establishment of a cap-and-trade system for carbon has been advanced before and failed, and no one can be sure a bill will pass anytime soon. However, the possibility of change invites a host of questions about how a carbon market would operate and what should be the rules.

---

<sup>\*</sup> MIT Sloan School of Management, MIT Center for Energy and Environmental Policy Research and the MIT Joint Program on the Science and Policy of Global Change, E40-435, 77 Massachusetts Ave., Cambridge, MA 02139 USA, E-mail: [ellerman@mit.edu](mailto:ellerman@mit.edu) or [jparsons@mit.edu](mailto:jparsons@mit.edu)

<sup>\*\*</sup> MIT Technology & Policy Program and MIT Center for Energy and Environmental Policy Research, E40-435, 77 Massachusetts Ave., Cambridge, MA 02139 USA, E-mail: [stephanf@mit.edu](mailto:stephanf@mit.edu)

This research was funded in part by a generous grant from the Doris Duke Charitable Foundation.

Markets for pollution are not new to the U.S. The first Bush administration pioneered their use with the creation of the SO<sub>2</sub> market under the 1990 Clean Air Act Amendments. The SO<sub>2</sub> market has generally been viewed as a success in terms of keeping the cost of emission reduction low. In fact, during the Clinton administration, the US used the widespread acclaim of the SO<sub>2</sub> market to buttress its diplomacy in favor of incorporating market mechanisms into global agreements regulating greenhouse gas emissions. Ironically, this diplomacy succeeded in making market mechanisms an integral feature of the Kyoto Protocol which, under the second Bush administration, the US then declined to sign. The European Union took the next step and made a CO<sub>2</sub> market the centerpiece of its own strategy for regulating emissions. The European Union's Emissions Trading System (EU-ETS) is now by far the largest emissions market in the world – more than 20 times the size of the US SO<sub>2</sub> market – and the European CO<sub>2</sub> price is the global benchmark.

The design of a cap-and-trade system raises a myriad of fundamental choices that will be hotly debated: how stringent should the cap be, which sectors of the economy should be covered, how should the allowances be distributed, and what should be done with any revenues earned from the sale of allowances. In this paper we focus on one component of the cap-and-trade system: the markets that arise for trading allowances after they have been allocated or auctioned. The efficient functioning of the market is key to the success of cap-and-trade as a system. How well have the EU CO<sub>2</sub> market and the US SO<sub>2</sub> market functioned? Are there any important lessons for the better design of a US CO<sub>2</sub> market?

The cap-and-trade approach to regulation stands as an important alternative to the traditional command-and-control approach. And in its short history, cap-and-trade has had some major successes in achieving significant pollution reductions at low cost. Nevertheless, these days, perhaps more than ever, there is a great amount of mistrust in commodity markets in general and financial trading in commodities in particular. This mistrust is especially great among some of those people most fervently advocating action to reduce carbon emissions. The popular mistrust in markets can lead policy makers to try to minimize the role of markets and trading. But such an approach compromises the ability of a cap-and-trade system to work. After all, the principle underlying a cap-and-trade system is that the forces of the market should be harnessed and exploited to produce low cost emissions reductions. The efficient functioning of markets is useful. The more market flexibility that can be built into the system, the better. Of course, vigorous oversight and proper regulations to weed out and prevent abuses are essential ingredients of an efficient market. But a regulatory structure that inhibits trading and thins the market or that restricts flexibility in meeting the cap ultimately raises the overall cost of the system.

The recent history of both the European CO<sub>2</sub> market and the US SO<sub>2</sub> market can illustrate these points. In both cases, the cap-and-trade system has worked relatively well, and the value of trading and flexibility are clear. And in both cases, there have been problems that can be traced back to insufficient flexibility and obstacles to trading that could have been avoided with modest amendments to how the systems were designed. While these problems have been small, an accurate diagnosis is important to shaping the future design of a US system. Both the European CO<sub>2</sub> and the US SO<sub>2</sub> cap-and-trade

systems allow some amount of inter-temporal flexibility in meeting the cap, so that companies can bank already issued allowances for use in future years. But in each system, where that flexibility has been restricted – whether in law or in practice – the result has been an inefficient evolution of price and therefore a slightly higher cost to the system. And where that flexibility has been expanded, the result has improved the efficiency of the price process and lowered the cost to the system. Both systems allow unrestricted trade in allowances, including the creation of futures markets. However, the decision to allocate allowances freely to companies presumed to be the natural shorts immediately thins the size of the market that actually develops – in both a direct and an indirect way – and reduces the ability of companies to take advantage of the inter-temporal flexibility that is allowed. The lesson for the current debate in the US is that efforts to impose strict regulations and oversight are to be encouraged insofar as the purpose is to encourage active trading and an efficient market, but efforts to impose limits on trading and to minimize inter-temporal flexibility and the role of the marketplace should be opposed.

## **THE EUROPEAN CO<sub>2</sub> MARKET**

The EU-ETS is a classic cap-and-trade system. For the first phase, running from 2005 through 2007, emissions were initially capped at 2.1 billion tons CO<sub>2</sub> annually. The cap covered emissions from more than 10,000 installations in the 27 countries of the EU, encompassing electric power and industries such as pulp and paper, metals, refining and cement. Emissions from households and transportation were notably not a part of the system. Allowances for emissions equal to the total cap were distributed annually. Most allowances were allocated free of charge to affected installations, although a small

number were auctioned. Companies were then free to buy and sell allowances throughout the EU, so that a market in allowances arose. Each year a company would have to report the CO<sub>2</sub> emissions for each of its covered installations, and then surrender sufficient allowances to cover those emissions. A company with emissions greater than the number of allowances it had been allocated would have to cover the deficit by purchasing additional allowances. A company with emissions less than the number of allowances it had been allocated could sell the surplus.

The system is now in its second phase, which runs from 2008 through 2012. The level of the cap has been reduced to 1.9 billion tons annually – slightly less than a 10% reduction from the first period cap on a raw basis, not adjusting for changes in industries and installations covered by the cap. Air transportation will be included within the cap starting in 2012 with a corresponding adjustment to the cap, and steps are being taken to move to a markedly greater use of auctions for distributing allowances instead of free allocations. As always, there is ongoing debate about the right level for future caps, about which industries should receive any free allocations, and about links to other countries. Otherwise the system largely continues in the same form as originally designed.

## **Trading**

An active market has developed for allowances. There is both a spot market in already issued and valid allowances, and a futures market for allowance vintages not yet issued. Figure 1 shows the growing volume of transactions through time. Total volume in 2005 was just over 262 million tons CO<sub>2</sub>, which is a turnover of 0.12 when compared against the annual allocation. By 2008, total volume had grown to over 2.68 billion tons and a turnover of 1.41. In comparison, the CFTC Commissioner's quote at the top of this

article assumes a turnover of 10 times the assumed annual allocation in the US, so it is clear that the \$2 trillion figure is not likely to be reached in the first few years of a US system unless it is somehow markedly different from the European system. Much of the trading is done OTC through brokers and on electronic systems, but there are also several European exchanges. The largest exchange volume is in the futures contracts offered by the European Climate Exchange (ECX) through the Intercontinental Exchange (ICE) platform. Futures contracts are offered for maturities running out a number of years. Options are also traded. The NYMEX is now attempting to establish a foothold in the carbon market, although its presence is currently negligible.

As with other cap-and-trade systems, banking of allowances from one year into future years is allowed. This increases the flexibility of the system as a whole. Inter-annual fluctuations in the demand for emissions can be smoothed by a corresponding fluctuation in the allocation of supply across years to minimize the annual variation in the marginal cost of abatement. This lowers the total cost through time of reducing emissions to meet the cap.

The EU-ETS, however, has two peculiarities with regard to the banking of allowances. First, it created multi-year compliance periods with fixed endpoints. Allowances issued for the first period could be banked within the period, but could not be carried forward into the second period. This created a clear seam between the two periods with a discrete difference in the price at the close of 2007 and the opening of 2008.

Second, as an artifact of the details of how firms were to comply with the regulations, it was possible for companies to borrow allowances from the next year's allocation in order to cover the prior year's emissions. Although this was not an



intentional design feature, this borrowing increased flexibility and helped to smooth price fluctuations, but it too could only be done within a given compliance period.

### **Price Evolution**

Figure 2 shows the evolution of the CO<sub>2</sub> price in the EU-ETS through the first phase and into the first year of the second phase. We show the spot price for a first phase allowance, good for emissions in the years 2005-2007. We also show the futures price for a 2008 allowance, deliverable in December 2008. There are four things to note in the evolution of these prices.

The first period price began slightly below €10/ton CO<sub>2</sub>, which is close to what many analysts predicted at the time. However, the price then rose persistently to a range of €20-€30/ton into April 2006, a range that was widely considered too high.

A common explanation for the persistently high price during this early period has to do with the lack of the physical supply of allowances coming to the market. The initial allocation of allowances displayed an obvious pattern: the various industry sectors were fully allocated allowances equal to anticipated emissions, while the power industry's allocation was reduced by the anticipated reduction in emissions for the system as a whole. It was the power industry, now short allowances, that was best prepared for the introduction of the system. Because of the large size of many of the players in the power industry, because of their active engagement in other commodity markets such as electricity and fuels where similar trading was occurring, and because of the greater significance that a carbon price would have on their variable costs, the power industry was quicker to understand how a cap-and-trade system would work. On the other hand, many of the industrial players were smaller companies with less familiarity with trading

and how a cap-and-trade system would function. Thus, while the power industry which was short allowances was carefully hedging its carbon exposure and buying allowances to match anticipated power production on a short forward looking basis, many of the industrial players who were long allowances, did not initially come forward and offer their supply into the market.

Also contributing to the shortage of supply in this early stage was the delay in the establishment of allowance registries in several of the east European countries that were new members of the EU. Although the EU-ETS is established by the Union as a whole, and although allowances trade freely across the EU, the implementation of the system is done at the national level. Much of the anticipated surplus of supply originated in these east European countries, and the delay in establishing the registries meant that the supply could not come to the market.

There is some evidence that intermediaries stepped in to speculate against this high price, but the shortage of the physical supply at this early stage of the new market affected as well the ability of intermediaries to borrow a supply with which to execute an arbitrage. Because the initial allocation of allowances was made specifically to those installations expected to be emitting, only a marginal fraction of the total physical supply of allowances ever needed to be brought to the marketplace to begin with, thinning the total level of the market. Had a larger fraction of the allocations been made through auctions, then the physical market would have been thicker and the opportunities for intermediaries to operate would have been greater.

The second notable event in the price chart is a discrete price drop in late April 2006, when the price fell from €30/ton to around €15/ton. This followed the release of the

verified emissions data for 2005, which indicated a markedly lower level of emissions than had originally been anticipated and therefore a lower marginal cost of meeting the cap than had originally been expected. Arguably it would be wise to develop a more frequent reporting of inventories, to lessen the discrete impact of each single information release and to lessen the store of unreleased information that can be leaked at any date. But the annual reporting is most likely to remain for a regulatory system that is in such an infant stage of development.

The third notable feature of the price path is the gradual drop in the spot price for the first phase until it is almost zero throughout the final three quarters of 2007. This led to many ill-informed statements that the European system had been “overallocated” allowances and that the EU-ETS was a failure in reducing carbon emissions. The zero price is not a reflection of the allocation. Instead, it reflects the seam between 2007 and 2008 built into the EU-ETS’s use of discrete phases without any banking or borrowing allowed between the phases. The cap remained what it had always been, and aggregate emissions were below the cap due to some combination of error in estimating baseline emissions, abatement and the randomness of actual emissions. So, in the first phase, the EU-ETS succeeded in capping emissions exactly where it had started out to cap emissions, and there was no failure from the perspective of the original system’s goal. However, the seam between the two phases of the EU-ETS built into the price evolution a peculiar dynamic as the close of 2007 approached. The market as a whole faced a binary outcome at the close of 2007. If the market as a whole had emissions more than the remaining allowances, then companies that were short would have to pay a penalty equal to the price of a 2008 allowance plus €40. On the other hand, if the market as a whole had

emissions less than the remaining allowances, then the unused allowances were worthless. As the year marched on, the final price could only take on one of these two values. It was impossible to take on a final value in-between. As it happened, the final outcome became relatively clear early on, and the price fell close to zero.

The fourth notable feature is the futures price for a 2008 allowance, which ranged in the neighborhood of €15-€25/ton CO<sub>2</sub> throughout 2007, despite the collapse of the spot price for a first phase allowance. This separation between the prices for first and second phase allowances shows the effect of the peculiar seam that is a feature of the EU-ETS. The positive price for a 2008 allowance also shows that, despite the collapse of the spot price for a first phase allowance, corporate investment and operating decision with a horizon longer than a few months would still have to take into account the cost of carbon. The separation between the prices for first and second phase allowances also illustrates the virtue of permitting banking and borrowing of allowances so as to smooth the marginal cost of abatement across years and avoid exactly this sort of price development.

### **The Spot-Future Relationship and the Term Structure of Futures Prices**

Within each phase of the EU-ETS, futures prices for delivery of allowances at different dates exhibit a very simple relationship. Since a company is indifferent between which vintage of allowance they deliver to cover their emissions, the only reason to pay a different price for delivery of an allowance in December 2006 than for a delivery of an allowance in December 2007, for example, is the time value of money. Therefore, with the exception of the early part of the first phase, when price quotes for futures did not necessarily reflect actual traded prices and the physical market as a whole was illiquid,

spot and futures prices of different maturities within each phase have moved closely together, with the basis being an interest rate.

In other types of commodities, such as agricultural and fuels, futures prices of different maturities often diverge, reflecting the different impact of short-term supply bottlenecks and short-term demand pressures and the consequent varying marginal cost of storage and marginal convenience yield. This differential pricing of different maturities is a useful tool for optimizing resource allocation through time in the face of varying cost and other factors through time. But for carbon, the social cost of a ton of emissions is always the same, regardless of the year of emission. Therefore, designing the terms of trade accordingly, so that the prices of emissions at different dates are distinguished only by the time value of money makes sense. In this type of system, a change in the expected marginal cost of abatement – whether because of changing technologies, changing economic growth, or an anticipated tightening of future caps – will affect the level of the spot price and the level of the full term structure of futures prices, but it will not affect the shape of the term structure as so often happens with other commodities.

## **THE UNITED STATES SO<sub>2</sub> MARKET**

Prior to the creation of the EU-ETS, the US SO<sub>2</sub> market was the premier example of a successful cap-and-trade program. Created under the Acid Rain Program of Title IV of the 1990 Clean Air Act Amendments, trading in SO<sub>2</sub> allowances began in 1995. The program is widely regarded as a success in cutting emissions at low cost. Early estimates placed the cost of cutting emissions and therefore the price at \$500/ton SO<sub>2</sub> or more. However, from its inception in 1995 through year-end 2003 the market price of an

allowance never exceeded \$220/ton. Although the US SO<sub>2</sub> market has been overshadowed by the EU-ETS's CO<sub>2</sub> market, the SO<sub>2</sub> market recently experienced an amazing price spike that exposed problems with the design of the market and that offers important lessons for the design of a possible US CO<sub>2</sub> market.

Under the Acid Rain Program, SO<sub>2</sub> emissions have been capped – since 2000 at 8.95 million tons, 10 million tons below 1980 levels. The majority of allowances were allocated on a grandfathering principle: existing emitting installations obtained annual allocations for 30 years according to a specified formula. A small number of allowances are auctioned each year. Allowance allocations are maintained in electronic accounts in an EPA database. Each year companies must report their emissions, and then surrender allowances equal to those emissions. Allowances are vintaged by year. Companies may bank allowances to cover emissions in future years. But allowances cannot be used to cover emissions in a year prior to the vintage of the allowance – i.e., there is no borrowing. Allowances can be freely bought and sold, with all transactions in the actual allowances being recorded in accounts on the EPA database.

For most of its life, the US SO<sub>2</sub> market has been a brokered market with relatively low volume in a small number of transactions. All transactions are recorded in the EPA's database. Total turnover of allowances between economically distinct organizations (as opposed to transfers between entities owned by the same parent) reached a peak of nearly 13 million tons in 2001 and equaled 10 million tons in 2005. Thus, turnover slightly exceeded the annual allocation of allowances, which is a low ratio compared to most commodity markets and comparable to the level reached in the EU-ETS in 2008. Early attempts to create exchange-traded futures contracts failed, but the effort was recently

renewed with some success. At the close of 2004, the Chicago Climate Exchange (CCX) launched a futures contract, followed in mid-2005 by the NYMEX. Options became listed in 2007. Volume in the CCX contract, the more active of the two, was negligible in 2005. In 2006, it had climbed to 723,000 tons and in 2007 risen to more than 9 million tons, roughly doubling the level of trading in the underlying physical market. Mid-year results for 2008 show volume continuing to increase markedly.

Initially, the ability to bank allowances across years worked as expected. The program had been introduced in two phases, with the first phase (1995-1999) being less stringent than the second phase (2000 and later). Companies chose to cut emissions more than required during the first phase, accumulating a large bank of allowances. Then, when the second phase discretely lowered the allowed level of emissions, companies drew on the balance of allowances in the bank, producing a “smooth landing” to the lower emissions cap in the second phase. This smoothed the marginal cost of compliance across the two phases, which was reflected in the smooth price for allowances across the two phases, quite unlike the seam that characterized the transition between the first two phases of the EU-ETS.

### **A Price Spike**

There still exists a sizeable bank of allowances. Nevertheless, in 2005-2006, the SO<sub>2</sub> market experienced a startling spike in the price – see Figure 3. During 2004, the price grew from less than \$220/ton up to slightly more than \$700/ton. During the first three quarters of 2005 the price continued its marked climb, passing \$980/ton. Then, in the last two months of 2005 it rose another \$600/ton, reaching its peak on December 2 at \$1,625/ton. With the turn of the year the price began a nearly equally precipitous drop,

falling almost to \$600/ton by May. For the balance of 2007, the price fluctuated between \$400 and \$600/ton. In 2008, the price began another decline, including a precipitous drop in July to nearly \$130/ton. It ended the year close to \$200/ton. This unprecedented volatility demands an explanation.

Two fundamental factors are often discussed as possible culprits in creating this spike. The first is the implementation of the EPA's 2005 Clean Air Interstate Rule (CAIR). Prospectively, CAIR imposed a stricter cap on the SO<sub>2</sub> market, first in 2010, and then stricter still in 2015. Clearly CAIR would raise the cost of an allowance. With banking allowed, anticipation of this should increase the current price of an allowance, leading to a smooth transition to the lower cap. Indeed, much of the price rise in 2004 and early 2005, after the rule had been proposed and while it was under public discussion, reflects anticipation of this fundamental change in the stringency of the cap.

Nevertheless, CAIR is not a good explanation for the full spike. The December 2005 peak of \$1,625/ton far exceeded estimates made during that time which forecasted that the cost of meeting the stricter standards would not exceed \$600/ton of SO<sub>2</sub>. More to the point, the passage of CAIR cannot explain the key feature of a spike, which is a price run-up followed immediately by a price drop. CAIR could justify a price increase that persists through time, but not a short-term spike. The existence of a large bank of allowances should have smoothed the effect of any price rise, leading to a gradual approach towards a persistently higher price level, and not to a sudden run-up in prices followed immediately by a collapse.

Ultimately, in a surprise ruling in July 2008, the District of Columbia Appeals Court vacated CAIR in its entirety. Then in December, the Court reversed itself in part by



reinstating the rule on a temporary basis, pending resolution of the underlying objection motivating the Court's decision. Obviously, the Court's decision explains the price drop in 2008, some in rumored anticipation of the decision and then a final sudden drop upon the official release. But the Court case is entirely irrelevant to the spike in 2005-2006. There were no interim procedural events or other elements of the case behind those price movements, and the sweep of the Court's ruling was entirely unexpected at that time.

The second fundamental factor often discussed is a disruption in delivery of low-sulfur coal from the Powder River Basin (PRB) of Wyoming to power plants in the Midwest. Track failures struck both the Union Pacific and the Burlington Northern Santa Fe railroad in May and in October 2005 creating a bottleneck that significantly reduced deliveries. In addition, a pair of coal mines had extended outages. The price of low-sulfur coal trading in the Midwest peaked in December 2005 at a level triple the price a year earlier. The shortage in low-sulfur coal forced power companies such as Xcel, WE Energies, Entergy and Alliant Energy, to shift to higher sulfur coal with attendant higher SO<sub>2</sub> emissions. Consequently, the demand for allowances was suddenly higher, driving the price up.

The railroad disruption to delivery of low-sulfur coal, in itself, is also not a satisfactory explanation for the spike. While the disruption was real, the temporary nature of the disruption was well understood. While the spot price of coal for immediately delivery understandably spiked, an SO<sub>2</sub> allowance is a different commodity. There existed a large bank of valid allowances that could readily be drawn upon to smooth the price across the temporary disruption. The price level might have risen modestly in response to this unexpected draw on the bank, but there should not have been a further

doubling or even tripling of the price nor the subsequent sudden collapse back to the level justified by CAIR.

### **The Restricted Float**

Both the implementation of CAIR and the supply disruptions to PRB coal did raise the immediate demand for allowances by specific companies. But ultimately, these companies found themselves squeezed to pay a very high price during 2005, despite what would appear to have been a plentiful supply of banked allowances. The owners of the banked allowances – overwhelmingly other electric utilities – did not come forward to meet this demand and profit off of the spectacular rise in price far above its long-term level. Why not?

There are three features of the design of the US SO<sub>2</sub> market which together assure that the float – the number of allowances actually available for trading in the market – is very small.

First, by originally distributing allowances to natural shorts, the system exploits trading only for marginal adjustments across power plants and other shorts in response to evolving variation in the cost of abatement. Most of the distributed allowances held by the natural shorts will simply be kept in their accounts until they are eventually surrendered to cover emissions. This dramatically thins the size of the overall marketplace, reducing liquidity.

Second, the original allocations are made free of charge, and so are held on the books of the companies at zero tax basis. If a natural short ultimately uses the allowance, it recognizes a gain in the value of the allowance, but, by definition, the gain is offset by the realization of a loss in the form of an incurred liability to pay for the emissions. So

the timing of the realization of the tax asset naturally matches the timing of the realization of the pollution liability, leaving a neutral tax implication. However, if a natural short attempts to profit from a temporary price spike by selling a banked allowance, later covering its exposure by a repurchase at a lower price, the short effectively accelerates the realization of the tax asset but does not accelerate the realization of the tax liability. This acceleration of the taxable gain is a penalty that must be weighted against the profit earned from arbitraging the developing price spike.

Third, many of the companies that are natural shorts and that hold the banked allowances are regulated utilities. The regulatory rules, both explicit and implicit, mean that neither the shareholders nor the management may capture any profit from arbitraging a developing price spike. To execute the arbitrage, the company is “speculating” with its stock of emissions allowances. The regulatory body may view profits from such a speculation as something that should be passed along to customers in the form of lower rates. On the other hand, a failed speculation with the stock of emissions allowances may be viewed as imprudent gambling with the electricity customers’ assets, so that the management is effectively penalized on the downside. Consequently, allowance banks created from free allocations to regulated utilities cannot be expected to fulfill the price smoothing function to the fullest degree possible.

In part as a consequence of these three features, financial intermediaries have not played a large role in the marketplace, except as brokers. They have not generally held on their own account a significant stock of the allowance bank. A number of financial intermediaries did begin to step into the market to accumulate some of the allowances,

starting primarily around 2005 and in connection with the attempt to start a futures market.

This lack of float in the US SO<sub>2</sub> market set the stage for the price spike of 2005-2006. When underlying fundamentals such as the implementation of CAIR and the interruption of deliveries of low sulfur coal require a sudden increase in trading to reallocate the burden of abatement across installations, the thin market is unable to handle the pressure and the allowance price adjusts to reflect this short-run burden as opposed to the long-run equilibrium. The bank fails to adequately serve the mission expected of it.

## **CONCLUSIONS**

The histories of the European CO<sub>2</sub> market and the US SO<sub>2</sub> market demonstrate the value of inter-annual flexibility in the use of emissions allowances—i.e., the banking and borrowing of allowances over time. Banking smoothed the SO<sub>2</sub> price in the US across the transition from Phase 1 into Phase 2 in 1999-2000. Where banking has been restricted, as in the EU-ETS CO<sub>2</sub> market across the seam at 2007-2008, the prices in the two phases diverge sharply, and opportunities to lower costs by reallocating emission reductions through time are lost.

Even where banking is legally allowed, the efficacy of banking depends upon the liquidity of the market and other institutional design features that encourage or discourage companies from exploiting the opportunities. The 2005-2006 price spike in the US SO<sub>2</sub> market illustrates this problem. The free allocation of allowances tends to restrict the full realization of the value of banking. By its very nature, it reduces the need for trading and so thins the market. But also through the interaction with the tax and

regulatory system, it reduces the value of exploiting the flexibility afforded by banking. As a consequence, the price impact of short-run shocks to the system are magnified, and a sub-optimal allocation of emission reductions through time results, raising the cost of the system.

The free allocation of allowances in the EU-ETS has presumably also thinned the market over what it might otherwise have been. The only identifiable impact was in the early period of the first phase, in 2005, when the market was imbalanced. As the EU-ETS is now committed to auctioning more allowances over time, we may see liquidity increase in the market.


While banking an allowance for use in a future period is a widely accepted and valued feature of cap-and-trade systems, borrowing from a later year's allocation to cover an earlier year's emissions has often been controversial. However, the initial period of the EU-ETS illustrates the potential value that borrowing has. As noted earlier, the high initial price in 2005 reflected the one-sided nature of the market at that time, with shorts in the market attempting to cover their exposure, but longs not yet bringing forward their supply. The ability to borrow from the 2006 allocation in order to meet the compliance obligation in 2005 limited the consequences of this initial start-up problem.

Although the concept of permitting borrowing is taboo for some, the same benefits can be captured by structuring the annual allocations of allowances appropriately. The goal of the system is to maintain a cumulative cap over a long period of time, without regard to the specific years in which emissions are made. The practice has been to translate the cap into annual allocations by simply dividing by the number of years in each phase so that each year's allocation is the same. In the current discussion of

a prospective US CO<sub>2</sub> cap-and-trade system, a declining cap is achieved through a sequence of step wise lower allocations. Typically these annual allocations are decided upon without the structure of banking and borrowing in mind, and then permission for banking – but not borrowing – is overlaid on top. This leads to the potential danger that the inability to borrow may be a binding constraint on the system in certain years. An alternative would be to settle on an aggregate cap, and then choose a front-loaded path of annual allocations which sums to this aggregate cap, but which also assures that the system builds up an adequate bank in the early years. This nullifies the importance of a constraint against borrowing and assures the system of the benefits of flexibility. A front-loading of allowance allocations would also minimize the effect of start-up problems like the imbalanced market that afflicted the EU-ETS in 2005. The ability to shape the time profile of the annual allocations – constrained by the aggregate cap – is a degree of freedom in the design of a cap-and-trade system that has been overlooked in the discussion.


If Commissioner Chilton's forecast of a \$2 trillion futures market in carbon is to come to pass, it will take not only time and patience, but also careful attention to the design of the cap-and-trade framework to assure the requisite flexibility in trading as well as the efficient functioning of the market. Ultimately, a flexible and efficiently functioning market serves the purpose of lowering the cost of reducing carbon emissions and advances the goals for which the cap-and-trade system is intended.

Figure 1  
Monthly Volume in EU-ETS Allowances, Spot + All Futures



Source: Point Carbon.


Figure 2  
EU-ETS CO<sub>2</sub> Price History



Source: Point Carbon and ECX.



Figure 3  
US SO<sub>2</sub> Price History



Source: Cantor-Fitzgerald. Data is the monthly index. Therefore, daily prices within the month, e.g., at the peak of the spike, may exceed the monthly number shown.