Interim report of Complex Systems Approach to Controlling Carbon Emissions

Zhenyuan Zhao & Neil F. Johnson

Physics Department, University of Miami,

FL 33126, U.S.A.

(Dated: August 4, 2009)

Figure 1 shows a schematic describing the carbon market structure. Emissions in the market are capped at a limit L over a period of T time steps (taken to be days). The average cap for emissions per day is thus $\bar{L} = L/T$. The objective is for the companies to emit exactly L units of carbon pollutants each month. If this limit is exceeded then the amount of carbon emitted into the atmosphere is too high. On the other hand, if the aggregated emissions are too low then this suggests some wasted production capacity. The only information given to the companies after each day is whether or not the emissions level was above or below the cap.

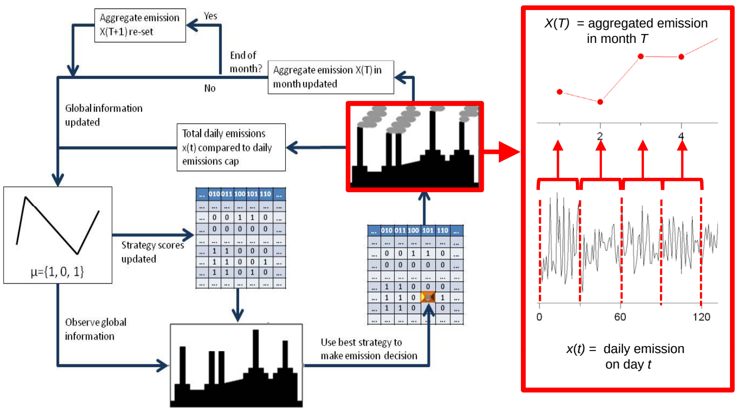

FIG. 1: Schematic diagram of the carbon market model and the resulting time series.

The performance of the carbon market is assessed through an analysis of the time-series for the emissions, see the right hand side of Figure 1, specifically the mean and maximum of the aggregated emissions over a fixed period of time and the standard deviation (volatility σ) about the mean. The results in this article show that within the basic constraints of the model, companies are able to organise themselves to hit their collective monthly emissions target, with minimal fluctuations in the aggregated emissions each month in the absence of any external regulator controlling the market. Moveover, government control does somewhat improve the hit with respect to the mean emission, but also introducing a bigger fluctuations (volatility) by disturbing the market. Also there may exist an "optimal" region of memory, the amount of previous days considered useful for prediction, that the effect of external control on the daily and monthly behaviours is almost invisible, which means the system can respond quickly to the outside. These features provide some insight into the behaviour of a feasible carbon market, which would be helpful for regulator as well as participants when it comes to the real market.

Here the market consists of N companies deciding whether or not to emit a unit of carbon on each day. Let the total number of companies choosing to emit on the first day be $x(1)$ ($x(t)$ is defined as the daily emission on day t). If $x(1) \leq \bar{L}$ the outcome, which is the only information that is made available to all of the companies, is the signal that there might be wasted emission capacity. If $x(1) > \bar{L}$ then the outcome is the signal that the emissions cap has been exceeded. In this case, a fine either economically or socially could be introduced for all companies emitting when the limit is exceeded. The outcome can thus be represented by a string of zeros (representing possible wasted capacity) and ones (representing over emission). One can then define $X(n)$ as the aggregated emission in month n. After day one, $x(1)$ is added to $X(1)$. Then on day two $x(2)$ is compared to $\bar{L}$ to determine the next bit in the outcome string and the aggregated emission is updated by adding $x(2)$ into $X(1)$. This process is repeated up to day T. After each period T the aggregated emissions is recorded and the aggregation process restarted for the next period.

In this market the companies do not communicate amongst themselves but interact through their common knowledge of the past history of whether or not the cap was exceeded, and the fact that each companies' decisions on whether or not to emit are influencing others' decisions through the possibility of emitting when many others do and thus exceeding the cap. All companies are assumed to have the same emission capabilities and to have the potential to emit, or not emit, one unit of carbon each time step. Each company holds a number of strategies that formulate decisions on whether or not to emit based on the m most recent outcomes of the market (also refer as memory m). Since there are two possible outcomes (cap exceeded or not exceeded) and two possible actions (emit or not emit), there are 2^{2^m} strategies in the strategy space. This number increases rapidly as companies are allowed longer memories. Before the companies begin competing in the market they randomly select s strategies from the strategy space with repetitions allowed during the assignment. Each strategy holds an action (emit or not emit) corresponding to each of the possible history strings and each company will use their best performing strategy to make each emission decision. The performance of each strategy is assessed through a scoring system **ss**, with the points assigned as follows. After each outcome is announced, each company assigns a point to each of their strategies that would have made the correct decision, and assigns a negative point to the strategies that would have made the incorrect decision. The correct decisions are emitting (not emitting) when the cap is not exceeded (exceeded). The incorrect decisions are emitting (not emitting) when the cap is exceeded (not exceeded). The best performing strategy each day is the one that has accumulated the highest number of points and this is the strategy that each company uses[1].

When $x(t) > \bar{L}$,

$$\mathbf{ss} = \begin{cases} 1 & \text{if not emit} \\ -1 & \text{if emit} \end{cases}$$

When $x(t) \leq \bar{L}$,

$$\mathbf{ss} = \begin{cases} -1 & \text{if not emit} \\ 1 & \text{if emit} \end{cases}$$

The market is thus composed of a number of companies competing to emit each day. Each company's decision to emit or not is determined by the best performing of a selection of strategies that they hold. The particular emission action each day is in turn determined by the recent emission history of the market. In this way there is feedback in the market. The companies interact through their contribution to the total emissions each day and the effect this has on the history of outcomes.

We have performed numerical simulations for markets with a range of emission limits L, time periods T and strategies per company s (see SI for more detail). The results provided here focus on a market composed of $N = 100$ companies, each possessing $s = 6$

strategies. The latter value ensures that companies will have a reasonable choice when selecting a strategy. The emissions are aggregated over a time period of $T = 30$, which can be considered to correspond to a month of 30 days. The emission limit for a month was fixed at $L = 1800$ units of carbon, corresponding to an average daily limit of $\bar{L} = 60$. On average, over all of the strategies, held by all of the agents, there should be an equal number of actions suggesting emitting as not emitting. Since $\bar{L} > N/2$, then in order to reach the desired monthly emission cap, the companies will need to emit more often than not, and thus the value of $L = 1800$ introduces a proclivity for companies to emit. Within many industries, emission is associated with production, and thus this desire to emit is a reasonable and realistic feature of market participants.

The values we are going to analyze (mean, maximum, standard deviation) provide three criteria by which the performance of the market can be judged. The optimal volume of carbon emissions each month is taken to be exactly equal to the emissions limit. This strikes a balance between minimising the volume of pollutants entering the atmosphere and maximising the level of industrial production. Similarly, the optimal maximum monthly emissions volume is also equal to the limit, thus never exceeded. However, neither the mean nor maximum aggregated monthly emissions are likely to ever exactly equal the limit and so these criteria can be more realistically stated as: in an efficiently performing market, the mean and maximum aggregated monthly emissions will be as close as possible to the monthly limit. The final performance criteria is that the emissions volatility is low. A low volatility is defined to be one which is less than the corresponding volatility for a 'random' market in which companies decide whether or not to emit each day by tossing a coin (either a traditional coin or a biased coin which elects to emit with a probability $p = \bar{L}/N$). This ensures that one does not have a situation in which the mean monthly emission appears close to the limit L, but the actual emission volumes fluctuate between well below the limit and well above the limit.

We also introduce some external intervention for the purpose of comparison, mimicking somewhat government control, where government reduces or increases the emissions capacity $L(n+1)$ for month $n+1$, by the amount that the aggregated emissions $X(n)$ in month n was above or below $L(n)$, which we called a "managed" system. As opposite to this, there is an "unmanaged" system, where the market is completely free with no external control, i.e., constant cap. Figure 2 shows some properties for both systems of the time-

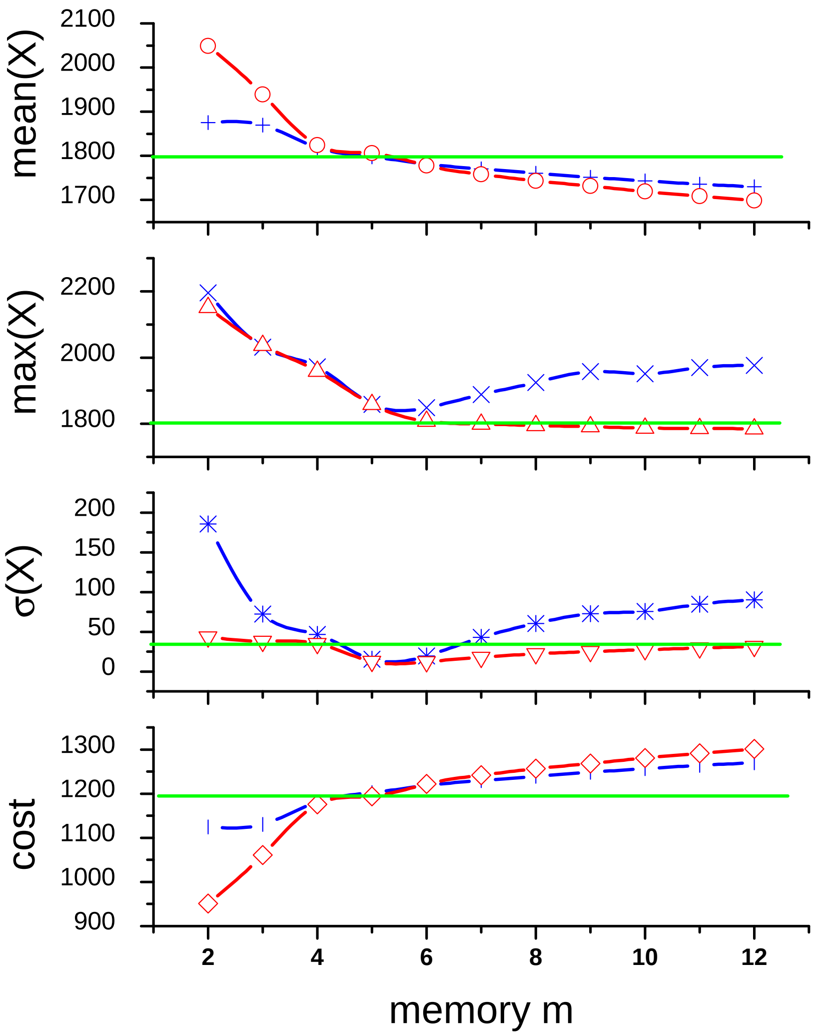

FIG. 2: Monthly emission measurements for $N = 100$, $s = 6$ from up to down as mean, max, volatility, and cost. Red is for an unmanaged system, while blue is for a managed system. It seems the management does pull the monthly mean emission closer to the designation, but at the cost of higher volatility.

series for the aggregated monthly emissions $X(n)$. The mean monthly emission X decreased monotonously as memory m increases for both systems, but the control from government does pull the mean towards the capacity limit, at a cost of higher volatility as well as bigger maximum for the high memory region. However, the reasons of the higher volatility due to the government control for the low and high memory regions are different. For small memory, i.e., $m \leq 4$, there are limited strategies available, the scoring system for them recycles quickly, because period $2 \cdot 2^m \preceq T$, which means there exists autocorrelation among daily emission x within a month which diminishes the volatility when looking it monthly (reduces $\sigma(X)$), see Figure 3(c) inset. But when the monthly intervention from the government comes in, it (autocorrelation) is destroyed to some extent. On the other hand, for high memory region $m \geq 7$, the intervention is in the way of the self-origination of the free market, since now period $2 \cdot 2^m > T$. In between, for $m \sim 5$ or 6, the mean and maximum monthly emissions are closest to the monthly limit for both systems. Based on the criteria above, this corresponds to an 'optimal' value for m. Importantly, at this range of m, the minimum in the volatility is significantly below the random coin toss limit for both managed and unmanaged systems, which means the model can adjust quickly to the government management, implying that "proper" choice of memory m might be helpful to a better performance irrespective to the outside environment.

Panel (d) of Figure 2 provides the overall cost to the governing entity as a result of a simple compensation scheme that compensates the companies that do not emit, in other word, a simple one-unit payout to any company not emitting on a given day. The choice of emitting or not-emitting is essentially cost-neutral to a give company. However for public relations reasons, and because they want to stay active in business, each company competes to emit on an 'undercrowded' day, or not-emit on an 'overcrowded' day. In this scenario, the payoffs become like El Farol, though we stress that the relevant time-series and management are generalized, aggregated versions.

Intuitively, there are two limits, one is a "non-learning" system, where the system has no feedback from the emission capacity, and does not have a tendency to hit the limit, the other one is a "perfect-learning" system, where there is strong feedback from the emission capacity, and the system reaches the limit perfectly well. In Figure 3, subplots (a) and (b) show the comparison of the results for daily emission. For the respect of mean emission (subplot (a)), the model is a good "learner", but when looking at the volatility (subplot (b)), it seems out of the picture, especially for low memory region, where herding/crowd phenomena is significant. Figure 3 (c) provides an illustration as the consequence of the different measurements of volatility $\sigma(X)$ (monthly) and $\sigma(x)$ (daily), see insert, where daily volatility has been scaled into the monthly level. Deviations between the two will be important for the practical pricing of the options contracts based on emissions, using standard derivative pricing theory. Here the price of the European call option depending on

Black-Scholes equation[2], is showed, which means price of option depends on the volatility σ, and then on memory m. So when the market starts behaving abnormally, for one reason or another, the memory, i.e., the amount of the previous days that are considered useful information, shrinks $m \to 0$. And the mispricing gets worse. Consequently, the by-effect of government control, bigger volatility, would also serve as a disturbance of the option market.

Lots of work has been done on stylized facts of markets, but most of them focus on price. In Figure 4, we look at the volume (V) of emission, taking it as an indicator of the number of companies who needed to emit, bought a permit, and emitted. Even if a fraction did this (say f) and hence $(1-f)$ were speculators, this would not change the quantity $(V-\bar{V})/\sigma$, since all quantities get multiplied by f which then cancel out. Our model captures the broad, apparent bimodal form of the real market, suggesting of something realistic at the microscopic level. But it also shows smaller occurrence of extreme event (our model gives zero when $abs((V-\bar{V})/\sigma) > 2$), which means the proposed competition mechanism might provide better control of the fluctuations.

Our analysis suggests a potential alternative structure for a carbon market, which requires minimal global control. Companies compete with each other for emission capacity and do not communicate directly to make their emission decisions, but instead infer the best action from global market information. Despite this limited communication, the companies are able to organise themselves to consistently reach their collective emissions target. In a broader context, the model suggests that companies are able to respond the changes in market conditions with different speed depending on the length of memory. The external control seems helpful to the respect of mean monthly emission, but inducing a much bigger volatility which makes the market very unstable. Moreover, the model captures some feature of the real EU Carbon market data, and might suggests a better control of the fluctuations through the collective competition.

---

[1] Tied best-performing strategies are broken by random choices.

[2] Black-Scholes equation: $V(x,t) = x\Phi[d_1] - X_s e^{-r(t_0-t)}\Phi[d_2]$, where $\Phi[z] = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{z} e^{-0.5y^2}$, $d_1 = \frac{ln(x/X_s)+(r+0.5\sigma^2)(t_0-t)}{\sigma\sqrt{t_0-t}}$, $d_1 = \frac{ln(x/X_s)+(r-0.5\sigma^2)(t_0-t)}{\sigma\sqrt{t_0-t}}$

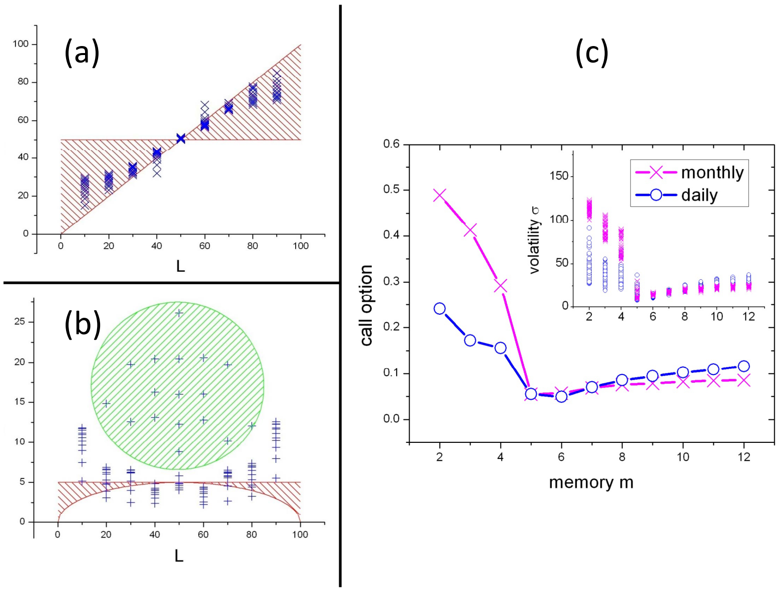

FIG. 3: The daily emission measurements for $s = 6$, and for (a) and (b) m changes from 2 to 12. (a) the daily mean $< x >$ as the change of L. The red shade is the region between no learning/adaption (horizontal boundary) and perfect learning (slope equals one). (b) the same set as (a) except looking at the daily volatility σ. The green shade shows high volatility due to the "crowd effect", while the red shade is the region between no learning (horizontal boundary) and perfect learning (convex curve). (c) Comparison of the European call option price resulting from different measurements of volatility according to the standard derivative pricing theory, Black-Scholes equation, where in this figure, risk-free interest rate $r = 0$, current value x is set to the individual mean $< mean > /100$, strike price X_s equals $\bar{L}/100$, for respect of individual company, the volatility σ is also scaled by dividing $1/\sqrt{100}$ into individual company's level, and the time $t_0 - t$ is set to 1, one day before expiration. Blue curve is the results from monthly measurement, while purple one is from daily measurement. The insert is the anomalous scaling of volatility of emissions, from daily to monthly scales. The purple crosses are the standard deviation calculated by taking the daily volatility and multiplying by the square-root of 30, which would be exactly the monthly volatility if the time-series followed a random walk.

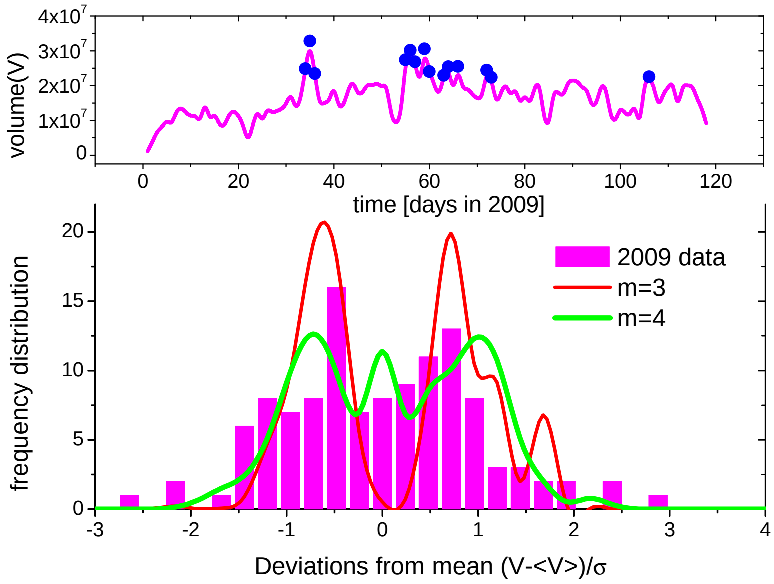

FIG. 4: The upper penal is the daily time series of the first 118 trading days' total volume in 2009 from ECX EUA Futures Contract, in tonnes of CO2 (EU Allowances). The lower penal is the frequency distribution of $(V - \bar{V}/\sigma)$ comparing with our model of memory $m = 3$ and 4.