GRISTMILL A blogful of leafy green commentary

Carbon policy details: Part 1

Carbon policy is close to getting the macro right, but plenty of smaller decisions remain

Posted by Sean Casten at 12:00 PM on 26 Mar 2008

My recent exchange with Gar has made it clear that there is a wide gulf between those details of carbon policy that are theoretically optimal and those which actually impact carbon reductions. Or, to be blunt, those that come up in our weekly staff meetings as actually affecting our decision to consider potential carbon reduction projects and those which simply elicit groans around the conference room of the "great intent, why did they screw up the execution?" variety.*

From our perspective, the good news is that our policy does finally appear to be moving not only toward putting a price on CO2 emissions, but getting the really important details (like auction vs. allocation) right. The bad news is that most of the other details are still wrong.

But those details are, well, complicated. And I don't know how to describe them succinctly. So consider this post just a teaser, with more to follow. All of the political candidates are now making the right noises about auctioning (instead of allocating) pollution allowances. And while that's far from settled, I'm not going to add anything to that discussion that hasn't already been Gristed. However, there is still a wide gap between the other details of what is conventionally believed to be good GHG policy and those that will quickly and cheaply lower GHG emissions. Specifically:

1. If "additionality" matters
2. Whether we let markets or government set the price of GHG emissions
3. Whether carbon contracts are denominated as single point "spot" or long term "strip" contracts

I will follow up with posts on each of these, but in the meantime suggest that there is a very simple way to evaluate any issue related to carbon policy, including the three noted above: **Will they encourage investment in capital projects to reduce carbon emissions?** I'll close with a suggested approach that is way simpler and more effective than any of the options currently being considered.

More to come ...

*Background detail: Our company has raised $1.5 billion to invest in projects that will satisfy our mission of profitable greenhouse gas reduction. Needless to say, this gives us a pretty strong incentive to think about how potential carbon reductions will impact our ability to invest that money. Thus the reference to the cheer/groan ratio around our staff meeting tables as an indicator of carbon policy effectiveness.

Carbon policy details: Part 2

Does additionality matter?

Posted by Sean Casten at 4:15 AM on 27 Mar 2008
The first follow-up to my recent post on carbon policy details.

First, a note to non-carbon-wonks: "Additionality" is a term of art in the world of carbon policy. It describes the degree to which a given activity causes *additional* carbon reductions -- the idea being that we shouldn't pay for carbon reductions that were going to occur anyway. As a fantastic oversimplification, suppose your car broke down and you had to ride your bike to work. The principle of additionality says you shouldn't be paid for the carbon you didn't emit. (You would have ridden anyway -- what choice did you have?) But if there's an increment of money that would tip you over into getting *rid* of your car and *always* riding your bike, that's additional.

Theoretically, great idea. Practically? Stupid.

To understand why, go back to the test I posited in my earlier post: **Does the metric increase or decrease the rate at which we invest capital to lower GHG emissions?**

The answer for additionality is not what you'd expect, for rather subtle reasons.

First off, let's note a couple truths:

1. No one's wallet is infinite.
2. Any reduction in GHG emissions is going to require a change to the status quo and, therefore, some capital investment.

Note that the capital investment may be small (in the case of a CFL) or large (in the case of CCS). It may represent a really good, no-brainer investment (CFLs, energy efficiency, etc.) or a really bad investment (again, see CCS). But someone's got to make a decision to spend that money. And since all wallets are finite, **any policy that preferentially directs resources towards lower-return investments is a policy that fails to maximize GHG reduction**.

Let's make it concrete. Suppose I have a million bucks. Should I invest that million bucks in something that saves 1,000 tons of GHG emissions per year and saves me $500,000/year in energy costs, or should I invest it in something that saves 100 tons of GHG emissions per year and costs me $500,000/year in operating costs with no associated savings? That should not be a hard question ... yet additionality tests make it so.

Why? Because additionality is a *qualitative test*. Too often it ends up being boiled down to financial metrics. Good investments, by virtue of being good, are judged to fall into the "you would have done that anyway" box, while bad investments, by virtue of being unable to attract sufficient capital, are deemed worthy of public incentive. Implicit is that public funds should only be put toward shoddy investments!

Worse, these tests are inadvertently hostile toward energy efficiency, which is the low-hanging fruit in the GHG-reduction orchard. If I can build a power plant that is 10 percent more efficient than the grid, it will generate a little bit of profit, but probably not enough to justify the capital investment. That means the resulting CO2 reductions could almost certainly be justified as "additional," since I can't make the investment without getting a public incentive to do so.

Now suppose instead I want to build a power plant that is 50 percent more efficient than the grid. That one has much lower operating costs and is much more profitable ... which makes it more likely to fail additionality tests.

And so **the additionality test will drive capital toward suboptimal investments and slow the rate at which we lower atmospheric GHG concentrations**.

In practice, this makes carbon markets that rely on additionality tests extremely transaction-cost intensive. Potential project investors are at pains to explain and document why they're not going to do this project without public support. Some of those explanations are true, some are bogus. But in all cases, it's a pain in the butt to do all that justification (read: lots of time and money spent to get paid for carbon reductions). That means fewer carbon reductions are going to made. Investors just don't want to go through the hassle of all that explainin'.

The irony is that we have good examples of carbon reduction markets that don't work this way. Renewable energy credits are assigned as an attribute of the power, on a $/MWh basis. Make more renewable MWh and you get more $, rather than having to spend more time explaining why you deserve payment *despite the fact* that you're making money.

We need to do the same with carbon emissions. Let's stop worrying so much about whether or not investors are making too much money lowering carbon. Far more important is that we start lowering carbon.

Carbon policy details: Part 3

Carbon taxes vs. carbon trading

Posted by Sean Casten at 10:30 AM on 28 Mar 2008

This is the third post in a series about details we are still getting wrong in the climate policy discussion. See also part one and part two.

There is no shortage of economic analysis and policy discourse that shows that carbon tax and cap-and-trade methodologies can deliver economically equivalent outcomes. The general consensus -- at least today -- seems to be that since they're equivalent, it really comes down to politics, and it's politically difficult to do anything with the word "tax" in it, so we'll do cap-and-trade. I like the conclusion, but the rationale is pure bunkum.

To understand why, we need only go back to my simple test of any climate policy proposal: **the degree to which it encourages investment in capital that lowers atmospheric greenhouse gas concentrations.**

Cap-and-trade and carbon taxes do not pass the test equally.

A carbon tax provides no direct revenue to carbon reducers

Suppose you're me. You've got investors who want you to invest in projects that reduce GHG emissions. They also want to invest their money as rapidly as possible, and to earn as much money as possible from those investments. All of that makes them very keen to figure out how different carbon policies affect their investment thesis.

You now find yourself in a board meeting, trying to explain to them how you will realize financial value from a carbon tax, given your expertise identifying and building projects that reduce GHG emissions. Let's walk through your conversation:

Q: Will your projects get paid for their beneficial GHG impacts?

A: No. The tax gets charged to emitters. Since my projects don't emit carbon, I don't have to pay the tax. But I don't get any of the revenue.

Q: That still gives you a cost advantage relative to your competition, right?

A: Not precisely. In the long run, the tax could increase the cost of electricity as produced by more carbon-intensive suppliers.

Q: "Long run"? "Could"?

A: Yes. Utility costs don't immediately translate into higher rates. First they have to go through rate cases, and there is a time lag between when their cost structure goes up, when they file for a new rate case, and when that rate case gets approved.

Q: But you can at least assume it will have a demonstrable impact on the competing price of power, right?

A: Not really. I don't know how the rate case will apportion rates across different rate classes, so I don't know whether my customers will be affected in a predictable way by the carbon tax. It is possible that they will simply impose those costs on other rate payers.

There is also a move afoot in the environmental community that would prohibit utilities passing the costs of greenhouse-gas abatement along to their rate payers; in that case, the impact of the tax would be to lower the profit margins of regulated utilities, but it would not have any impact on the competing price of electricity.

Q: But surely you can structure power contracts with your hosts such that they give you some upside from the resulting carbon tax, should it come, right?

A: Maybe, but that's a hard commercial sell. Would you pay me money today on a gamble about the direction that Congress will take on tax treatment for carried interests to private equity funds? Probably not -- and for the same reasons, it's unlikely that my customers would give me upside on a tax bet.

Q: So what value should we place on carbon reductions for your projects, given this carbon tax?

A: Zero.

This (Socratically) is the crux of the problem with carbon taxes. **It is a stick upon the emitter, without any direct carrot for the reducer.** And the financial value to the reducer is therefore only realized depending upon the manner in which the emitter's prices change -- but this is far from precise. (Witness all the manufacturers whose profits get squeezed when fuel costs rise because they have no direct way to pass those costs along to consumers.) By contrast, a system that allows emitters and reducers to *trade* allows for direct financial benefits to those who do the right thing -- and a benefit that can be negotiated, built into contracts, and used to affect investment decisions.

Tax policy is a blunt tool

So suppose we could fix all the above. Maybe we do a balancing tax and use tax increases on one side to offset tax reductions on the other side. Does that solve the problem?

Directionally, it's an improvement. But it is still a poor substitute for a market-based trading system, for the simple reason that tax policy is an indirect tool.

Tax incentives are a great vehicle, *if you are a taxpayer*. Individuals pay income tax, and mortgage interest tax deductions are a nice perk. But people who invest in energy projects are not taxpayers, at least not in the near term. Why? Because of depreciation and debt. If I spend a lot of capital on a project, I get depreciation shields to lower my taxable income in the early years of project operation. I will also almost certainly debt-finance some portion of the project, and my interest payments will also provide a tax shield. For most energy-related investments (e.g., all investments that either emit or avoid the emission of CO_2), you can count on not having any taxable income for the first 7 to 10 years of the project's operating life. And if you're not paying a tax, a reduction in your taxes isn't worth much.

So what do you do? You sell off your tax liabilities to someone with a "tax appetite," in the jargon of the trade (read: someone with lots of profits that they'd prefer not to pay tax on -- typically, banks). They will pay you if you assign them your rights to those tax losses, based on some discounted cash flow stream, giving you an immediate cash flow -- but essentially giving away some portion of the gain to the tax equity buyer (after all, they're not going to buy all those losses from you without getting some benefit).

Don't get me wrong -- this transaction is not a bad thing. Indeed, this is how most wind projects get financed. But the end result is that of all the dollars the government allocates from the treasury for The Good Thing, less than 100 percent of those dollars flow to the person who does The Good Thing. Which means that less than 100 percent of the investments that could theoretically be incentivized by this tax policy actually get incentivized, while the balance essentially goes to bank fees.

So how do you fix this? Two potential ways: One, make the federal payment a revenue payment rather than a tax offset. This is what Section 451 of the recently passed energy bill does, and it's huge. The net cost to the feds is the same, but for every dollar paid by the treasury, there is a dollar going to incentivize projects. Two, allow bilateral trades between those who cause the pain and those who provide the gain.

Both of those put the incentives in the right place and create real incentives for those considering investments in carbon-reducing technologies. And neither are amenable to a tax-driven approach.

Carbon policy details: Part 4

Spots vs. strips

Posted by Sean Casten at 2:29 PM on 31 Mar 2008

This is the fourth post in five-part series on the details required to get carbon policy right. See also parts one, two, and three.

We now get into an issue that will seem a bit arcane, because no one's talking about it, at least not explicitly. But it's a real choice, and in many conversations about carbon policy we are implicitly getting it wrong.

Should we price carbon in spots, or strips? Or, to take it out of financial jargon, should we:

1. set up markets such that people who are selling or buying emissions credits have to go to the market with each incremental ton to determine what the price will be (a "spot" market), or
2. set up markets such that buyers and sellers can enter into long-term contracts for the emissions they will produce/reduce (a "strip" market)?

Before talking about carbon, we need to take a brief foray into electricity deregulation, specifically back to the early 1990s, when academics were developing rules for how a deregulated market would work, based on broader theories of how markets work. As a reminder, **the central test of a good carbon policy is whether or not it encourages investment in carbon-reducing technologies**. Replace the word "carbon" with "electricity" and you've got a good test for whether the attempt at electric deregulation worked. It didn't.

In those states which elected to deregulate their wholesale power markets, the driving theory was that if you provide a *spot* price for power, and allow any buyer or seller to trade at that spot price, you will liberate the power of markets.

While there are many complicated details to these transactions, that is basically still the way all these markets work, whether at PJM, ISO-New England, NYISO, or any number of foreign markets. There is a fair amount of regular trading volume on these markets today, and so it may sound odd to hear someone say they don't work. But they don't, for the simple reason that it's really hard to get comfortable spending billions of dollars of capital based on a gamble about where future prices are going to go.*

A thought experiment may be helpful. Imagine that you have a billion dollars -- about enough to build a modestly sized coal plant (or, if you prefer, around 400 MW worth of wind turbines). You have the good fortune to be planning your project in a deregulated state, so you don't need to get a bunch of permission slips from the local utility to build your plant. You just need to get permits, connect to the transmission system, and build. Now, for the sake of argument, we will stipulate that the average price on those power markets for the last year has been sufficiently high that -- if your plant was operating last year, and selling into those markets -- you would have earned $200 million/year in net revenue. That sounds pretty good, right? Five year payback, or an 18 percent return on your investment over the next 15 years. Would you build the plant?

This pretty quickly raises an obvious question: what's the price on the market going to be next year? The answer: **You have no idea.** Sure, you can predict. You can hire consultants to make forward price forecasts. But when all is said and done, if you build this plant, you are placing a bet with risks you cannot control. What if the economy slumps, demand for power falls, and the spot price collapses? You lose. What if spot prices go up? You win. Which is more likely? You don't know.

This is not to say that no one will build the plant. Some people have a higher tolerance for risk and will build anyway, in the belief that they have some underlying knowledge of where power markets are going. Some of them will be right. Some of them will be terribly wrong (witness Calpine's stunning $16 billion bankruptcy which, at core, was the result of a bad bet on the way that gas and electric prices were going to move on those spot markets).

But here's the rub: we didn't have to design the market that way. Indeed, there is no real market that works that way. (And by "real," I mean a market that traces its origins back to one dude who had something another dude wanted and negotiated a price, as opposed to ones whose rules were crafted by academics and regulators.) If you live in New England and buy fuel oil to heat your house, you have a whole choice of payment options. You can pay each month based on the price of oil. You can prepay and lock in your price for the year. You can pay a little more on delivery in exchange for a predictable fixed contract. And so on. Ditto for any number of other volatile commodities that we all buy on a regular basis.

And, ironically, the old, regulated utilities never build plants on spot prices. When they build a plant, they hold a rate case and then lock in their price. (They will also then buy long-term "strips" on their fuel contracts to ensure that they don't get pinched between future fuel and electric price volatility.)

Note that this is not to say that strips are always better than spots -- simply that if you want people to invest capital, you need to provide the *option* to sell on a long-term strip, short-term spot, and any crazy hybrid of the two. And those options naturally emerge whenever governments simply allow buyers and sellers to meet up, negotiate deals, and get out of the way. Risk-averse buyers will naturally gravitate towards risk-tolerant sellers and vice versa, ultimately creating a mélange of spots, strips, futures, swaps, hedges, derivatives, and all those exotic-sounding increments that are the hallmark of a functioning market. (For those not familiar with the terms, don't worry -- they can all be thought of as different ways to bundle risks, with higher risk "flavors" offering potentially greater -- but more volatile -- returns and lower risk options offering stabler -- but generally lower -- returns.)

Relevance to carbon markets

You can probably see where this is going. How are we going to price carbon?

As a tax? Notwithstanding my prior post, this approach isn't even as good as spot, since it is a price set from a regulator on high, stipulating a fixed price with no market correction to capture the vagaries of supply and demand.

As an auction, with periodic re-auctions to reset the price, as Gar suggested? That gives us something like a spot but has all the problems that we've seen in the last decade in the electric sector.

As a government-established trading floor, modeled on the electricity sector? That simply repeats the problems we've already faced in the electric sector.

Here's the salient point. As of today, the only way you can buy or sell carbon (in the U.S.) in voluntary markets is on a spot basis. And the majority of the trading structures that are being considered in Congress are implicitly spot markets. And spot markets will have little impact on the decision-making process of those who want to invest money in projects to reduce greenhouse-gas emissions.

To the extent that there is a simple solution, it is this: **trust markets**. Let buyers and sellers meet up and trade in whatever fashion they like, and you'll get what you need. (With suitable government oversight, of course.) And if regulators cannot trust markets, then at least take advantage of the government's balance sheet to build in long-term contracts from the get-go. But don't assume that a spot price alone is sufficient.

*Point of candor: In the clearing prices for power (PJM, ISO-NE, etc.) that were set up as purely spot markets, one can now find other parties -- primarily financial traders -- who will buy and sell long-dated contracts, finally giving some semblance of "strips" to these markets. However, those transactions remain external to the regulated market, and it bears noting that it has taken about a decade for these players to arrive.

Carbon policy details: Part 5

The solution: Output-based standards

Posted by Sean Casten at 9:53 AM on 03 Apr 2008

This is the fifth and final post in a series on the details required to get carbon policy right. See also parts one, two, three, and four.

So far, I've done a lot of complaining -- which, in and of itself, is just, well ... whiny. Here, then, is a solution.

First, a very brief review:

1. A test of good carbon policy is whether it encourages the private sector to invest capital in projects that will reduce GHG emissions.
2. "Additionality" confuses carbon policy, by preferentially shifting investment toward less economic GHG-reduction technologies.
3. Carbon taxes provide sticks without carrots, and thereby provide no direct incentive to those who might otherwise use carbon pricing to invest in projects that lower GHG emissions.
4. Long-term carbon pricing is necessary to encourage private sector investment. Spots alone will not.
5. Although not covered in this series, it bears repeating that auctions trump allocation.

Unfortunately, virtually all of the GHG-reduction strategies currently in existence (e.g., Kyoto, RGGI) or being contemplated (e.g., Lieberman-Warner, California AB 32) fail one or more of the prior tests. Moreover, all those actual/proposed bills are really complicated, with many moving parts that are rife for gaming -- or, more charitably, significant legislative error. Here, then, is a better approach: **output-based GHG regulation.**

Output-based standards: The simple version

The concept of output-based standards comes from the world of criteria pollutant regulation. In conventional environmental rules, pollutants are regulated on a so-called "input basis" -- which is to say, the more fuel you consume, the more pollution you're allowed to produce.

As illustration, consider a typical parts-per-million (ppm) standard. Let's say that you are a regulated source (boiler, power plant, etc.) with a permit that allows you to produce no more than 15 ppm NOx. You build your plant, put the appropriate pollution controls on, and check in at 14 ppm. Now watch what happens if you try to increase the efficiency of that plant. The denominator of your permit (the "millions") is a function of only two variables: fuel combustion and air stoichiometry (the amount of air consumed per unit of fuel). For most technologies, air stoichiometry is just about fixed, so your only real variable is fuel combustion -- which means that the more fuel you consume, the more total pollution you are allowed to produce, given a constant ppm.

Now, let's assume that you suddenly find a way to dramatically increase the efficiency of your process, thereby cutting your fuel use in half. Since stoichiometry is fixed, you also cut your air flow in half. Let's also assume that you cut your NOx emissions, but "only" by 40 percent. That's good, right? Lower fuel costs, less fuel combustion, lower NOx -- you would think so. But since the denominator (the millions) has fallen faster than the numerator (the "parts") you suddenly find yourself in violation of your permit. 14 ppm NOx in the efficient device has suddenly become 14 x 0.5 / 0.4 = 17.5 ppm NOx, and you can no longer run your plant. The result? You simply keep running your inefficient device, over-emitting NOx, over-burning fuel, and over-paying for energy.

This is a really dumb policy outcome. It's also ubiquitous. What would be vastly better is to go to an output-based standard, where your pollution is a function of the useful output of your project (e.g., lbs/MWh). Drive up your efficiency, and it's easier to comply with your permit. Drive down your emissions, and it's also easier to comply with your permit.

Remarkably, virtually all jurisdictions still use input-based standards, although a few (like Texas) have shifted to an output basis. Now let's look at carbon.

Output-based GHG standards

As typically framed, carbon emissions are regulated on an absolute basis (tons per year, or some variant thereof). This would seem to be a step in the right direction, but it's not. Our goal is to reduce global GHG emissions, not local emissions. And regulation that sets local caps has a really hard time setting an appropriate value for anything beyond battery-limits emissions. Witness the ongoing debate on biofuels on Grist -- if I replace my fuel oil boiler with a biomass boiler, do I get credit for eliminating 100 percent of my carbon emissions, or should I first have to pay for the CO_2 associated with fertilizer production, soil carbon depletion, and non-sustainable harvesting?

This is a legitimately difficult question. And it's hard on the positive side as well. Suppose I replace my boiler with an on-site cogen plant that generates all my heat and power. I no longer buy power from the grid, and I'm making power more efficiently than the power I used to buy, but locally, my emissions increase, and so I have to try to figure out how much credit I should get (in terms of CO2 displacement) for the power I'm no longer buying. Again, hard issues.

Under current carbon rules, these dilemmas are addressed through some sort of offset rule (raising additionality issues, as Gar has pointed out) or on rather abstract theories that, since carbon prices will eventually affect those other fuels, you'll eventually realize some incentive. Which is sort of like promising to buy someone a birthday present once they turn 80.

All of these problems go away if we shift to an output-based standard.

Nationally, our greenhouse gases come from three sources: electricity generation (about 40 percent of U.S. GHG emissions), transportation (27 percent), and thermal energy generation (33 percent). (The latter category includes not only residential and commercial space heating, but also a wealth of industrial operations -- drying, calcining, melting, etc.) Let's take electricity first.

In 2006, the U.S. produced 2,393 million metric tons of CO2 emissions in the process of generating electricity, and we generated just over 4,000,000 thousand MWh of electricity. By division, that works out to about 0.6 metric tons per MWh. That means that anyone who builds a power plant with CO2 emissions exceeding 0.6 metric tons per MWh is making our life harder, and anyone who's beating it is making our life easier.

Under an output-based approach, we simply set a standard for all power plants at 0.6 MT/MWh and mandate that anyone who's over that number has to pay for any excess pollution, while anyone who is under gets paid for their good deeds. Who do they pay? Each other. Measurement is easy, since both parties have fuel bills and electric meters -- so, from a regulatory perspective, we simply require that they submit audited records of both at the end of the year, along with evidence that they bought or sold to get to 0.6 as appropriate.

To make sure that this structure caps CO2 emissions (e.g., to ensure that it doesn't increase as total MWh production increases), we simply reset the cap every year based on actual CO2 emissions. If total CO2 emissions have increased by 2 percent, we reset the cap in the following year to 0.6 x (1 - 0.02) = 0.588. (A variant of this structure is that you can recalculate the MT/MWh factor every year based on actual data. As the payment streams incentivize low-carbon generation -- and discourage high-carbon generation -- this continual re-averaging would have the effect of steadily lowering the allowable level of pollution -- a cute trick that requires no regulatory approval to tighten emissions standards.)

Now watch what happens. The structure has immediately created both carrots and sticks. If you install a solar panel (0 MT/MWh), you have 0.6 MT/MWh to sell. (And if you can drive up the efficiency of your solar panel, you get more tons to sell.) If you build a coal (1 MT/MWh) plant, you've got to buy 0.4 MT/MWh. Additionality? No need -- just good guys and bad guys, selling or buying GHG emissions. Spots vs. strips? Solved -- after all, the coal plant wants a long-term fixed price just as much as the solar plant does -- they simply have to agree on a price. Goal- vs. path-based regulation? Done. The price is on carbon, not a technology. If a sexy new yet-to-be-dreamt-of technology comes along, it gets the same access to carbon pricing as everyone else.

The same mechanism can be applied to thermal energy generation, although the math is a bit trickier. (And indeed, to accurately capture the value of cogeneration, you have to factor in the value of thermal energy.)

The benefits are that the government need only set up the rules and then provide the appropriate oversight. It can then get out of the way and let market forces figure out how to optimally price and deploy technology. In doing so, it will naturally drive capital towards the most cost-effective carbon reduction approaches.

The tricky bits

That said, there are a couple of complexities:

1. The model doesn't deal with transportation very well. In theory, you could calculate tons/mile, or some equivalent metric. In practice, though, that would be really hard to monitor. Unless someone has a better idea, this probably means that you'd still have to impose some type of carbon tax on transportation fuel for economy-wide coverage. (Note, however, that the tax could be set to equal the carbon-content of a given fuel multiplied by last year's average carbon price in thermal and electric markets on a $/ton basis, providing a linkage between the two markets.)
2. For similar reasons, residential fuel use is hard. Again, it may be that a sector-specific tax is ideal.
3. Biomass would require some additional paper trail to quantify which fraction of the carbon associated with the biomass is truly renewable and which part is not.

I've got no way to fix these complexities -- but compared to the massive complexities that work into tax and/or cap-and-trade models, this is comparatively quite simple. It also has the benefit of being fiscally neutral. With the exception of the wrinkles for transportation and residential (which could be eliminated), the buyers and sellers balance exactly, so there is no net economic cost -- simply a wealth transfer from polluters to cleaners. This means that we get big reductions in GHG emissions with no economic pain (and, better still, we see no net increase in energy costs, since every increase in the cost of a dirty MWh is offset exactly by a reduction in the cost of a clean MWh.)

Eager to hear comments from the Grist community. But if I were king, I'd roll this out tomorrow.