Energy Security Initiative
at BROOKINGS

COST CONTAINMENT FOR CAP-AND-TRADE:

DESIGNING EFFECTIVE COMPLIANCE FLEXIBLITY MECHANISMS

Bryan K. Mignone

SEPTEMBER 2009
Policy Brief 09-03

Energy Security Initiative
at BROOKINGS

COST CONTAINMENT FOR CAP-AND-TRADE: **DESIGNING EFFECTIVE COMPLIANCE FLEXIBILITY MECHANISMS**

Bryan K. Mignone

SEPTEMBER 2009
Policy Brief 09-03

About The Brookings Energy Security Initiative

The Energy Security Initiative (ESI) is a cross-program effort by the Brookings Institution designed to foster multidisciplinary research and dialogue on all aspects of energy security today. ESI recognizes that public and private choices related to energy production and use will shape the global economic, environmental and strategic landscape in profound ways and that achieving a more secure future will therefore require a determined effort to understand the likely consequences of these choices and their implications for sound policymaking. The ESI Policy Brief Series is intended to showcase serious and focused scholarship on topical issues in one or more of these broad research areas, with an emphasis on targeted policy recommendations.

Contact for the Energy Security Initiative:

Lea Rosenbohm
Project Manager
(202) 797-6248
lrosenbohm@brookings.edu

COST CONTAINMENT FOR CAP-AND-TRADE: DESIGNING EFFECTIVE COMPLIANCE FLEXIBILITY MECHANISMS

This paper is the second in a five paper series on US cap-and-trade design

Bryan K. Mignone

Bryan K. Mignone is a fellow and director of research of the Energy Security Initiative at the Brookings Institution. His research focuses on the science, economics and politics of global climate change, with an emphasis on the design and implementation of market-based climate policies.

The author gratefully acknowledges financial support from the Energy Foundation.

Introduction

A viable long-run solution to the global climate problem will require a sustained political effort to confront and ultimately resolve two top-tier policy issues. First, those with the greatest stake in the outcome—the major economies, leading carbon emitters and most vulnerable nations—must broadly agree on a consensus definition of "dangerous anthropogenic interference" with the climate system and a method for translating such impact thresholds into global temperature, concentration and emissions goals. Second, upon agreeing on such an objective, these stakeholders must then confront the more daunting challenge of equitably distributing the total abatement burden in order to arrive at a well-coordinated, technologically feasible and politically sustainable set of national emissions reduction pathways.

A related paper discusses the challenges associated with the determination of global and national emissions reduction trajectories (Mignone, 2009). Here we imagine that an appropriate US emissions reduction blueprint has already been selected from the space of available alternatives and focus more narrowly on the set of design considerations that could enhance the overall performance of the resulting regulatory program. We start from the premise that cap-and-trade will be the primary policy vehicle through which any proposed emissions reduction schedule will be realized.

In the larger taxonomy of regulatory solutions, cap-and-trade is the logical alternative to a carbon tax, which has had relatively fewer proponents in the US legislative context. While either instrument can be used to motivate emissions reductions, and while both fundamentally rely on price signals to alter future investment behavior, the choice of cap-and-trade reflects a judgment about the practical and political merits of directly regulating emissions quantities. Substantively, cap-and-trade provides greater certainty about the quantity of emissions abatement at the expense of some certainty in the resulting price of emissions allowances.

In light of this tradeoff, advocates of cap-and-trade must address concerns about price and cost uncertainty in order to secure broad political support for the effort and to build a policy framework that is both cost-effective and durable over the long periods of time demanded by the climate system. Near-term uncertainty in prices—or more specifically, price volatility—is an unavoidable feature of real commodity markets, but one that is likely to be problematic in the carbon abatement context only if it is extreme enough to jeopardize the clarity of the underlying investment price signal. Maintenance of this signal is, after all, the primary motivation for establishing a carbon market.

In this paper, we focus on a key element of the response to such price uncertainty, namely the suite of compliance flexibility mechanisms that could be incorporated into the fabric of policy itself. We suggest that carefully designed temporal flexibility instruments, such as banking and borrowing, combined with a limited centralized authority to make subtle market adjustments, could eliminate most price volatility resulting from short-term economic dislocations.[1] When it comes to longer-term uncertainty and the possibility that sustained high prices and costs will threaten the durability of the policy itself, we suggest that a carefully-designed upper bound on the carbon price could reduce these threats without materially increasing the risk to the overall environmental integrity of the program.

This paper proceeds as follows: Section 2 discusses key market performance objectives relevant to the design of policy. Section 3 considers firm-level temporal flexibility mechanisms, and Section 4 examines centralized adjustment mechanisms. Section 5 discusses more explicit measures to control prices and costs, and Section 6 summarizes the major issues and discusses the ways in which policies could be combined successfully in the context of a real carbon trading regime.

[1] Of course, there are likely to be other sources of volatility in the market. Price risk from such volatility could be addressed using other common hedging tools, like futures and options contracts. For a discussion of such hedging instruments and the implications for the development and regulation of the future carbon derivatives market, see Pirrong (2009).

Market Evaluation: Sharpening Design Objectives

Previous experiences with market-based environmental policies in the United States have demonstrated that the economic advantages of such approaches can be significant (e.g. Stavins, 2000). To fully appreciate these benefits in the carbon mitigation context, consider a (hypothetical) example of a non-market policy in which government regulators require utilities to install CO_2 "scrubbers" on all existing coal-fired power plants (assuming in this example that such devices are commercially available). For any given operator, this narrow "command-and-control" mandate may or may not turn out to be the cheapest route to the desired emissions abatement goal. Moreover, even if it is for some, a "one-size-fits-all" approach neglects the possibility that it may not be for all, meaning that the policy misses some important opportunities for additional cost reductions.

Next compare this policy to one in which regulators simply impose a set of legally-binding caps on annual emissions and then step out of the way. With such caps in place, regulated entities (those with a legal obligation to surrender allowances for their emissions) could satisfy their requirement in a number of different ways, including ways that might not be fully anticipated at the time of regulatory enactment. For example, a utility could decide for itself whether it would make more sense to install a new CO_2 scrubber, to build a new nuclear or wind plant or to pursue an entirely different strategy. A cap-based policy is therefore "technology-neutral" and promotes so-called *how-flexibility*, or the ability for firms to meet a particular regulatory objective *however* they wish, in pursuit of the lowest cost outcome.

This simple example suggests that compliance flexibility is an important driver of economic efficiency. Intuitively, a policy that endows individual actors with the flexibility to make bottom-up compliance decisions will outperform a top-down, one-size-fits-all approach, because any single compliance strategy is unlikely to be optimal for all regulated actors, particularly when the scope of the policy is very broad. As we shall see, the benefits of decentralized decision-making extend well beyond those offered by technology-neutrality and how-flexibility.

To explore other potential opportunities for enhanced economic efficiency, consider the one offered by allowance trading. Hypothetically, a firm that could meet its regulatory obligation for $20 per ton of CO_2 would benefit from selling allowances to a firm spending $30 per ton, while the second firm would similarly benefit from buying such allowances, assuming in each case that the trade were executed between $20 and $30 per ton. In effect, the buyer of allowances would under-comply relative to its own regulatory obligation, while the seller would over-comply by an equal amount. In this way, the overall policy objective, a specified economy-wide annual emissions reduction, would not be threatened. This type of compliance flexibility is often called *where-flexibility*, because it allows firms to pursue abatement *wherever* it is least expensive, again in pursuit of the lowest cost outcome.

Both how-flexibility and where-flexibility are essential, defining features of cap-and-trade, in the sense that no specification of this type of policy could exclude them (one that did would no longer be cap-and-trade, but something else). In fact, one might say that the two flexibility features discussed so far are captured, respectively, by the "cap" and "trade" elements of the term "cap-and-trade." In other words, a cap-and-trade system could be operationally defined as a mandatory program in which government regulators enforce an economy-wide compliance obligation without specifying either the exact location or the means by which that aggregate obligation is achieved.

Given the obvious advantages of how-flexibility and where-flexibility, it is natural to examine the opportunities provided by *when-flexibility*, or flexibility in the timing of abatement. The justification for this type of flexibility might seem questionable, since shifting abatement across time would imply altering the prescribed sequence of annual emissions reductions and ostensibly the goal of policy itself. However, merely altering the time path of abatement without altering the cumulative emissions reduction does not actually threaten the environmental outcome, because the ultimate climatic response depends on the cumulative emissions release over the lifetime of the policy, not on the specific annual sequence of emissions.[2] The justification for inclusion of when-flexibility is therefore provided by a physical feature of the climate system itself.

The link between compliance flexibility, broadly defined, and economic efficiency provides the essential motivation for the establishment of a carbon allowance market, because a market is the policy vehicle best suited to ration allowances (i.e. enforce the cap and promote how-flexibility), exchange allowances between regulated entities (promote where-flexibility) and distribute allowances across time (promote when-flexibility). If the establishment of an effective and efficient carbon market is the basic goal of policy, then questions about the optimal design of cap-and-trade are, more fundamentally, questions about the optimal design of a carbon allowance market. For this reason, in considering various cap-and-trade design possibilities, we will evaluate policy choices according to their ability to facilitate two broad market objectives: *price stability* and *cost viability*.

How do these two broad objectives relate to the three flexibility goals described above? As mentioned earlier, how-flexibility and where-flexibility are more or less guaranteed once a market is established, so these opportunities for cost reductions do not require further policy attention once the decision is made to utilize a market approach to abatement. On the other hand, temporal flexibility does require additional policy attention, since it is not automatic if the emissions reduction obligations are prescribed annually, as they tend to be for practical reasons. Attention to temporal flexibility is important, not only because it is the third pillar of economic efficiency described above, but also because it promotes the related market objective of price stability.

[2] This so-called stock pollutant assumption, which underlies much of the existing literature on carbon mitigation, would be best described as an approximation (see, e.g., Wigley et al., 1996). In reality, the global climatic outcome probably does depend somewhat on the time path of abatement (see, e.g., O'Neill and Oppenheimer, 2004), but this approximation is likely to be acceptable at the national level, because small deviations in the time path of abatement for one country will not materially alter the global abatement path.

To see the connection between temporal flexibility and price stability, consider the implications of an inflexible policy that prohibits deviations from the prescribed annual abatement path. In this case, firms would be forced to internalize cost shocks in the period in which they occurred, causing prices to fluctuate over time in response to uncontrollable and unpredictable external changes in economic growth, weather conditions, and so on. If firms were instead allowed to shift their obligations across time using banking and borrowing provisions (discussed at greater length below), the impacts of discrete shocks could be spread across multiple compliance periods, leading to a more stable trajectory of prices. One may conclude from this that a policy that maximizes temporal flexibility is tantamount to a policy that maximizes price stability or minimizes price volatility.[3] A smoother trajectory of prices, in turn, yields a more efficient outcome overall, as well as a more transparent price signal to market actors seeking to internalize the price of carbon in strategic investment decisions.

The second main objective introduced above—cost viability—does not follow directly from the inclusion of temporal flexibility, because a temporally efficient trajectory does not necessarily guarantee one that is aligned with the underlying political willingness-to-pay for climate mitigation. If costs exceed this political threshold, because (for example) low-carbon technology turns out to be considerably more expensive than initially anticipated, no amount of temporal flexibility will succeed in realigning actual prices with acceptable prices. So to the extent that sustained high costs provide an existential threat to the system, provisions to keep the market within the bounds of political enforceability need to be explicitly considered. Put differently, if the basic existence and continuity of the policy is in jeopardy, then efficiency is a second-order concern.

While cost viability is thus an important prerequisite to price stability, provisions designed to enhance price stability are politically easier to realize, because temporal flexibility lowers compliance costs without threatening the cumulative environmental integrity of the system. It is thus win-win, or at least win-not-lose. Cost viability is somewhat more challenging as a policy objective, because placing an explicit upper bound on costs (and therefore on prices) could lead to revision of the underlying emissions goals. Of course, failure to appreciate the constraints on political enforceability could also result in similar (or even more draconian) revision due to direct political intervention. The challenge in designing such provisions therefore derives from the need to balance the tradeoff between environmental integrity on the one hand and political enforceability on the other.

In the discussion that follows, we consider the set of policy instruments available to enhance the two broad objectives of market design. We start by discussing banking and borrowing as the simplest instruments by which temporal flexibility (and thus price stability and economic efficiency) could be realized. Since the successful operation of these provisions depends critically on firm-level foresight and long-term market durability, we next turn to policy instruments that could be used to address deficiencies in these conditions. In particular, we consider the ways in which limited centralized intervention could mitigate market dislocations unlikely to be resolved through the exercise of firm-level flexibility. We then consider mechanisms designed to address cost viability and to enhance the potential for long-term political durability, since existence of the system is a critical prerequisite to all other objectives, including price stability.

[3] As discussed earlier, minimizing volatility does not mean *eliminating* volatility, since volatility may result from other market factors as well, including basic informational uncertainty and, in some cases, from trading activity itself.

Market Efficiency: Enhancing Temporal Flexiblity

The previous section provided the basic justification for the inclusion of provisions that enhance when-flexibility. Here we confront the challenge of implementation, beginning with a simple banking example. Suppose the emissions reduction schedule codified in legislation reflects the optimal economic response to a prescribed cumulative abatement goal, given the information available about future abatement costs at the time of enactment. In the event that costs turn out to be lower in the early periods of the program, say because baseline emissions growth turns out to be lower, then the new optimal emissions path would exhibit more reductions early and fewer reductions later, relative to the prescribed targets.

This revised path is the one that would be realized if firms were allowed to bank (i.e. hold) allowances for future use and acted with foresight to realize the cheapest overall outcome. From a regulatory point of view, there is no compelling reason to restrict such behavior, because firms that bank allowances are simply over-complying in early years relative to the prescribed emissions reduction targets. In fact, not only will banking *not* threaten the environmental outcome of the program, it may actually *enhance* the regulatory objective in other ways, because firms that build up stocks of unused allowances add to their books tangible assets whose value is directly tied to the price of carbon. As a result, the strongest constituency in favor of the continuity of the system may actually emerge from regulated industry itself, providing additional political glue to bind the system together over the long periods of time required by the climate problem (c.f. McKibbin and Wilcoxen, 2008).

In theory, borrowing is the natural complement to banking. If costs in early years turn out to be higher than anticipated, then the revised optimal emissions path would exhibit fewer reductions early and more reductions later, relative to the initially prescribed targets. That is, firms would choose to borrow permits in the present and repay them in the future, shifting their obligations from periods of high cost to periods of lower cost. Although this kind of inter-temporal transaction is the mirror image of banking, the practical design implications are quite different, because borrowing (unlike banking) requires regulators to assume the obvious risks associated with lending.

In a world of perfect regulatory enforceability, early under-compliance resulting from borrowing would not threaten the system in any significant way. Firms that borrowed too many allowances early on might regret that decision later, but the regulator would still be able to enforce targets, and the cumulative emissions outcome would remain unchanged. In the real world, of course, the possibility of imperfect enforcement must be explicitly acknowledged and then effectively incorporated into the initial design of policy.

An important case to consider is the possibility that some firms carrying emissions debts might go out of business before fully repaying their "loans." This would leave regulators with the option of redistributing the abatement burden among other actors in the system or sacrificing the long-term integrity of the program. Neither route is particularly appealing, but then again, defaults due to bankruptcy would probably not account for a large share of total abatement.

A potentially more serious problem arises if a larger number of firms borrow extensively in early years, believing from the start that regulators will be unable to enforce stringent targets later on for political reasons. In this case, the very act of borrowing enhances the likelihood of this outcome, because greater borrowing in early periods requires more stringent and thus more politically uncomfortable reductions later. As the enforceability of future targets becomes less certain, firms might actually borrow more, further undermining the regulator's ability to enforce future targets, and so on, making initial concerns over enforceability a self-fulfilling prophecy.

Of course, problems with lending are not unique to the carbon market. Lending in any context comes with the unavoidable problem of default risk, which leads most lenders to place strong, legally-binding constraints on those to whom they lend, regardless of whether the debtor in question is an individual, a large corporation or a sovereign nation. The set of tools available to a sophisticated lender is broad and includes, among other things, strict quantity limits (i.e. an upper bound on how much will be lent), interest on borrowed principal, and collateral as a partial hedge against future default.

Each of these tools is broadly applicable to borrowing in the context of a carbon market. Two of them—quantity constraints and interest requirements—have been explicitly included in existing cap-and-trade proposals, including the Waxman-Markey bill, also known by its acronym ACESA. The version of that bill passed by the House limits the amount of borrowing to 15% of a firm's total emissions and requires 8% annual interest to be paid (in tons) on any borrowing, which must be repaid within five years.[4] No existing cap-and-trade proposals have thus far required collateral to be posted on borrowed allowances.

Of all the tools available, strict quantity limits are the most straight-forward but also the least refined, in the sense that the tradeoff between enhanced economic efficiency and enhanced protection against default risk is particularly stark. As the quantity constraint is relaxed in order to accommodate a wider variety of borrowing needs (i.e. protection against more severe cost shocks), the protection it offers against excessive borrowing and future default risk is diminished. On the other hand, as the quantity constraint is tightened to mitigate default risk, the ability to provide temporal flexibility is reduced.

The application of interest to borrowed allowances is a more refined mechanism for balancing temporal flexibility and protection against default risk. It works by effectively imposing a tangible penalty on borrowing. Firms that might engage in speculative borrowing—speculating, that is, on the future existence of the trading system—would reconsider that strategy if the costs associated with borrowing were sufficiently high. On the other hand, firms facing legitimately severe cost shocks would still be permitted to borrow relatively large amounts and would do so provided that the perceived efficiency benefits exceeded the implied borrowing costs.

[4] Text of this bill can be found at: <http://energycommerce.house.gov>.

The imposition of a concrete borrowing penalty thus helps to re-align firm-level incentives with the regulator's enforceable path, without artificially truncating the amount of flexibility allowed in more extreme cases. In fact, by more carefully considering the penalty required to bring about this realignment, the regulator could determine the optimal value of the interest rate to apply to borrowed allowances. In general, this parameter depends on assumptions about the discount rate that firms use to make compliance decisions, among other things.

While interest provides a more refined mechanism for trading off compliance flexibility and protection against default risk than explicit quantity constraints, this approach is not without drawbacks. First, note that interest is effective in mitigating the risk of default because the costs imposed on borrowing suppress incentives to overuse the mechanism. That very cost, however, also partially undermines the larger objective of cost containment, by making the economically efficient outcome more costly to achieve when borrowing is justified by market conditions. That is, the very feature that makes interest effective as a constraint is paradoxically what diminishes its effectiveness as a means of cost containment. Moreover, while the imposition of interest on borrowed allowances constrains excessive borrowing, it does not necessarily eliminate defaults altogether. When defaults do occur, this mechanism does not provide a means to compensate for lost abatement, a problem that leads us to consider the third type of borrowing constraint: collateral.

In a real allowance market, collateral requirements could be naturally implemented as a deposit on borrowed tons. In this case, a firm wanting to borrow an allowance would post a specified dollar amount (per ton) in advance, and the regulator would hold this amount in escrow for the duration of the loan, which it would later return (with interest at the risk-free rate) upon repayment of the emissions debt. For instance, a firm might borrow one ton from the regulator at $20 in the first compliance year and repay those emissions five years later, at which time it would receive $23 back from the regulator (arbitrarily assuming ~3% interest in this example).[5]

As a visible deterrent to overuse of the mechanism, the deposit closely resembles the interest rate constraint. A firm that was inclined to borrow in a speculative manner would reconsider that strategy if it were risking its own deposit on such a bet. Once again, by considering the penalty required to realign firm incentives with the regulator's enforceable path, one could estimate the value of the optimal deposit payment. Under the most straight-forward assumptions, the optimal deposit payment is equal to the *expected* allowance price.[6]

Despite the functional similarity between interest and collateral, the deposit avoids some of the problems that plague the interest rate mechanism. For example, in the event that defaults do occur, the regulator now holds collateral that can be used to pursue at least some of the missing abatement. In addition, the deposit is entirely refundable, so while firms betting on future enforceability face real costs, the ability to recover the efficient path when borrowing is justified is not threatened by the imposition of external costs (since the net present value of the combined deposit transaction—the payment and refund plus interest—is zero). More simply, the deposit mechanism lowers the transaction costs associated with borrowing, while

[5] For a deeper discussion of this concept, see Mignone (2009b). The interest rate paid on deposits could also be discounted further relative to the risk-free rate in order to ensure that firms do not use the borrowing instrument as an alternative investment vehicle.
[6] See Fell et al. (2009). The optimal deposit value also depends on assumptions about future price volatility.

still providing an effective deterrent to overuse, thereby increasing the odds of realizing an economically efficient outcome.

Ultimately, in a world in which firms operated with perfect foresight and in which the durability of the system were guaranteed, the individual incentives provided by banking and borrowing would be sufficient to bring about the temporally efficient outcome. However, since these two conditions (perfect foresight and durability) are not guaranteed in the real world, and since both are important prerequisites for the successful operation of the mechanisms discussed here, policies to support these additional conditions (or respond creatively to the lack of them) must be explicitly considered as part of the larger design effort. In the next section, we consider how centralized intervention might provide one such opportunity.

Market Stability: Centrally Adjusting the Allowance Supply

In the discussion above, we suggested that well-designed and carefully implemented banking and borrowing provisions could dramatically improve the efficiency of the carbon market by endowing firms with the flexibility to shift their compliance obligations across time in response to unpredictable cost shocks. We also noted that the success of these provisions rests on the foresight with which firms react to changing market conditions. In most circumstances, these idealized assumptions, while imperfect, are likely to be good enough to keep the market operating within reasonable bounds and close to its efficient level. At other times, however, firms may behave more myopically, with potentially significant implications for the efficiency, and even stability of the broader market.

These concerns are particularly relevant to the carbon market, because, as a novel regulatory construct, it may not be viewed as particularly durable by market participants. In light of these concerns, the remainder of this section discusses the regulator's role in actively managing the allowance market when firm-level actions fall short of what is required to maintain price stability. Enhancing stability, in turn, will improve overall confidence in the system and increase the amount of foresight applied to individual compliance decisions, thus reducing the need for intervention in the first place.

In periods of normal operation, decentralized firm-level decisions, like banking and borrowing, essentially act as negative feedbacks to small price perturbations, meaning that they tend to restore conditions that prevailed before the shocks were encountered. Consider, for example, the impacts of an upturn in prices. As prices move higher, firms will start to sell expensive allowances, borrowing them from future periods if necessary; since more permits will then be available for purchase in the current period, prices will begin to fall. In this way, the act of selling or borrowing allowances in response to an initial upside price shock will restore the system toward its steady-state level. Banking provides a similar negative feedback on the system when the initial price perturbation is downward rather than upward.

When the perturbations are larger or more sustained, however, and when firms do not act with sufficient foresight, the prevailing negative feedbacks, which act to promote stability, may give way to significant positive feedbacks that act to undermine stability. For example, firms that would be inclined to sell allowances under the conditions above might hold them under less certain conditions, expecting that prices will rise even higher in the short-run. If the upward trend itself encourages buying (rather than selling) for short-term gain, then firm-level decisions will enforce (rather than counteract) the initial perturbation, moving the system away from its efficient path. Put differently, as forward-looking firm-level behavior becomes increasingly myopic, individual decisions cease to restore the system toward its socially efficient level and may actually facilitate market instability.

While the appropriate policy response to this behavior is also debatable, regulatory attention to such issues seems justified in this context, because a carbon market (unlike most other commodity markets) is designed solely to provide a clear price signal to energy investors. In addition, as mentioned above, the carbon market is more likely to suffer from problems related to myopic behavior because, as a system invented by regulators, it may not appear to be very permanent. The bottom line is that the carbon market is both more susceptible and more sensitive (in terms of its performance) to myopic behavior than other "natural" commodity markets.

A hypothetical template for centralized intervention—dubbed a "carbon fed" by some—was proposed in the debate leading up to consideration of the Lieberman-Warner bill in the 110th Congress.[7] As envisioned by its proponents, this entity would be authorized to pursue several well-defined functions in the carbon allowance market, all intended to help manage prices in times of stress. In the original proposal, the set of possible actions included (a) the ability to expand constraints on borrowing (for example, by relaxing quantity constraints, lengthening the repayment period or lowering the interest rate), (b) the ability to relax quantity limits on offsets and (c) the ability to directly adjust the physical supply of allowances by essentially borrowing against future targets.

Each of these proposed mechanisms raises important questions about implementation. For example, it is not immediately clear *how* targets should be adjusted or *how* the offset supply should be expanded, even if one wanted to pursue these avenues. If the regulator makes frequent discretionary adjustments to defend a predetermined price, such abrasive intervention could undermine the performance of the market in other ways. For example, attempting to force adherence to a given price would discourage non-compliance entities from investing in the market and could also reduce firm-level banking and borrowing activities. In short, excessive discretionary control could drain market liquidity.

On the other hand, if centralized adjustments are made by applying specific, quantitative rules, the rules themselves run the risk of being arbitrary, disruptive or (quite likely) both. For example, while near-term adjustments to targets might move the market in the intended direction, successful implementation would require the regulator to first determine the threshold price at which to intervene and then quantify the nature of that intervention (in terms of the number of allowances to pull from future periods, etc.). In this case, the application of arbitrary rules triggered at equally arbitrary prices could cause disruptive shifts in the market, again undermining the very problem it was designed to mitigate.

The other major regulatory lever, the ability to expand the offset supply, would seem to suffer from many of the same problems. In addition, this approach begs an uncomfortable question: If more permanent, additional and verifiable (in short, high-quality) offsets are available at low cost, why should they not be allowed into the system from the start? Doing so would presumably lower the costs of compliance without jeopardizing the cumulative environmental integrity of the system. The most likely answer to this question is that

[7] See the paper titled "Cost Containment for the Carbon Market: A proposal," prepared by the Nicholas Institute at Duke University, July 2007. Available at: <http://www.nicholas.duke.edu/institute/carboncosts/ carboncosts.pdf>.

expansion of the offset supply would actually require one to sacrifice some emissions integrity.

Interestingly, once the tradeoff between the quantity and quality of offsets is explicitly acknowledged, a natural opportunity arises to consider the benefits of temporary carbon storage. Consider a somewhat extreme hypothetical example in which current allowance prices are $50, twice their long-run equilibrium value of $25, and in which two risky (non-permanent) offset tons are credited at 50%. This could be implemented by issuing one allowance to the submitting entity and one to the regulatory authority itself. The allowance injected directly into the market would slightly increase current supply, and the regulator could further increase supply by liquidating its own allowance for $50. The revenue generated from the latter transaction could be used to buy back two allowances in the market in a future period once the price dropped back to (or below) the equilibrium value of $25. This combined transaction would be both carbon neutral (the cumulative emissions reduction would not change) and revenue neutral (the net amount spent by the regulator would be zero).

The intuition provided by this specific example can be used to generate a more generally applicable "model rule" relating the discount rate on offsets to both the quality of the offset (its reversal risk) and current market conditions (the magnitude of the market price shock).[8] According to this rule, high-quality offsets—those operationally defined as permanent and additional according to well-defined certification criteria—would be credited one-for-one at all times. Under normal market conditions, lower quality offsets would be discounted according to their reversal risk, while in times of temporarily high prices, lower quality offsets would be discounted according to both their reversal risk *and* the economic value associated with transferring abatement from periods of high cost to periods of lower cost. That is, lower quality offsets would be discounted less heavily in times of high prices, effectively increasing the total supply of allowances and providing a stabilizing influence on the market. Such a rule could be applied subtly and continuously, without relying on arbitrary price thresholds, and for this reason, would avoid the most significant problems associated with other forms of centralized intervention.

In addition, greater attention to the offset supply would directly confront fears that lax certification procedures might flood the carbon market with cheap offsets of dubious quality, threatening both the investment price signal and the underlying environmental integrity of the program. Under the mechanism proposed here, regulators would be authorized to develop transparent but strict rules for certification, discounting all but the most high quality offsets. Under normal circumstances, this would suppress the offset supply sufficiently to prevent any erosion of the price signal or emissions integrity. In times of temporarily high prices, regulators could relax the discount rate applied to other classes of offsets, thereby enlarging the offset supply. This mechanism, while subtle, would greatly reduce the need to introduce other sorts of instruments, like a price floor, to hedge against the possibility of a market flooded by poor-quality offsets, because such credits would be allowed into the system only when prices were sufficiently high.

Ultimately, even the savviest regulators will not have enough information to determine the efficient price path in advance. This is the basic problem with discretionary intervention. Regulators can adjust the performance of the market in more

---

[8] For a detailed explanation of this methodology and a description of the specific rule, see Mignone, B. K., M. D. Hurteau, Y. Chen and B. Sohngen, "Carbon offsets, reversal risk and US climate policy," *Carbon Balance and Management* 4:3 (doi:10.1186/1750-0680-4-3). Available at: <http://www.cbmjournal.com/ content/pdf/1750-0680-4-3.pdf>.

subtle ways, however, to keep it within reasonable bounds that enhance foresight and encourage individual market actors to chart a more efficient course for themselves. The considerations above suggest that the best kind of intervention would utilize a transparent, economically-defensible rule that could be applied continuously over a range of market conditions rather than abruptly at discrete thresholds. In the next section, we turn finally to the most fundamental prerequisite for all of the mechanisms described so far: the durability of the market itself.

Market Durability: Ensuring Political Enforceability

The policy mechanisms discussed above are designed to facilitate a stable and efficient trajectory of carbon prices. Yet there is no guarantee, even if these provisions are implemented flawlessly, that the implied, fully efficient trajectory will be politically viable. The precise political willingness-to-pay for climate mitigation is obviously difficult to quantify, in large part because it will evolve during the legislative process, as policymakers gradually develop familiarity with the range of carbon prices associated with various quantity reduction proposals. If the prices implied by a particular set of quantity provisions do not attract sufficient political support during this process, then legislation based on such provisions is unlikely to succeed, meaning that it will need to be adjusted until the new implied price trajectory aligns with the underlying political willingness-to-pay.

These considerations suggest that the legislative process itself should accomplish most of what is necessary to produce a politically enforceable outcome, because close alignment between implied prices and willingness-to-pay will be required for passage. However, this argument neglects the role of economic and technological uncertainty in shaping the outcome and the possibility that the actual price trajectory could deviate markedly from the central case projections around which political support initially coalesced. Small deviations from the expected path would probably not cause major problems, but more extreme deviations could enhance the risk of direct political intervention. Put differently, a successful legislative result guarantees continued support for the most likely trajectory of allowance prices, but it does not necessarily guarantee support for *all possible* future trajectories of prices.

In response to this concern, a natural policy response would be to impose a price ceiling at the implied political willingness-to-pay. By introducing such a constraint in advance, architects of cap-and-trade policy would effectively preempt intervention by ensuring that market prices never exceeded the basic political enforceability threshold. In addition, they would enhance the credibility of the policy more generally, by explicitly acknowledging the range of prices that the regulator would be prepared to defend *ex post*. While conceptually quite simple, a price ceiling is somewhat more challenging to design, because two nontrivial questions must be resolved in advance. First, the precise value of the threshold (or a clear rule for deriving it) must be provided, and secondly, the specific operational details of the instrument must be carefully considered.

In approaching the first of these decisions, policymakers will need to make an informed judgment about the price at which future regulators or policymakers would be inclined to intervene. When it comes to implementation, the simplest way to maintain a given price ceiling would be to authorize the regulator to sell an unlimited number of allowances at that price, a policy construct often referred to as a *safety valve* (Jacoby and Ellerman, 2004; McKibbin and Wilcoxen, 2004; Pizer, 2002; Roberts and Spence, 1976). With such a mechanism in place, firms that could not meet their full compliance obligation at lower prices could purchase safety valve allowances in order to make up the difference. Despite its operational simplicity, the safety valve mechanism has been criticized on the grounds that it would, if triggered, relax the emissions cap and threaten the overall environmental outcome.

In a real trading system, the risk posed to the environmental integrity varies inversely with the value of the safety valve price. A sufficiently high threshold, whose chance of being triggered is perhaps 5%, is clearly less problematic from an environmental point of view than a lower threshold whose chance of being triggered is 50%.[9] However, if the chosen value of the threshold is not high enough to mitigate concerns about environmental integrity, then the trigger could also be coupled to conditional provisions that restore the cumulative emissions outcome over longer periods.

A few different design approaches are possible to promote this objective. First, regulators could require individual firms to perform additional future abatement in proportion to the number of safety valve allowances purchased, effectively duplicating the deposit mechanism discussed earlier. Of course, if the payback window were relatively short, this mechanism would not be particularly well suited to mitigating sustained periods of high cost, while if the window were much larger, it would be more difficult to implement at the firm level.[10]

Alternatively, the regulator could implement repayment at the *system level* by adjusting future system targets in response to early emissions overages. It is worth noting that this design would probably not restore emissions integrity if costs remained high and if safety valve allowances continued to be available for purchase in future compliance periods. On the other hand, if the safety valve were removed at some discrete time in the future, the very act of removal might create other problems.

A final solution would bring forward a fixed number of emissions allowances from future periods *before* the start of compliance and sell those allowances in a secondary or "reserve" auction in which the floor price would be set equal to the chosen trigger value.[11] This arrangement is yet another variation of system-level emissions repayment, in which repayment is guaranteed because the intertemporal transfer is effectively performed before trading begins. Despite offering a creative solution to the previous problem, this design comes with two other potential drawbacks. First, reducing the total number of allowances in the system by pulling them out initially will tend to *increase* allowance prices on average, thus partially undermining the broader cost containment objective. This concern can be mitigated by pulling the permits from periods farther in the future, which

---

[9] The average initial allowance price under ACESA projected by the US Energy Information Administration is about $18 per ton CO_2 (EIA, 2009). Thus one might assume that a price ceiling at this value would be triggered with ~50% probability, given inherent compliance cost uncertainty. In order to lower that probability to ~5%, the price ceiling would need to be much higher, perhaps two or three times greater than the average projected price, depending on the uncertainty in the compliance costs. And just as the expected allowance price rises at the interest rate, the price ceiling itself should rise at this rate from its initial value.

[10] Also note that if the safety valve payment were not returned upon repayment of the emissions debt (as under the deposit), then use of this mechanism would be quite costly to firms (c.f. Mignone, 2009b).

[11] For a deeper discussion of this mechanism, see Murray et al. (2008). Full text of this paper is available at: <http://www.nicholas.duke.edu/institute/wp-costemissions.pdf>.

are less likely to influence current prices. Second, since the "reserve pool" of allowances is fixed, the threshold price—and the political enforceability threshold—could be exceeded, meaning that the mechanism might fail to perform its core objective in some (rare) instances.

In light of these concerns, policymakers will need to choose an appropriate mechanism by carefully evaluating the tradeoffs. A basic safety valve would provide the ultimate backstop if future costs turned out to be much higher than anticipated, and if the value were chosen carefully, it would also ensure that the *expected* deviation of emissions from the prescribed total reduction would be negligible. However, it cannot *guarantee* that the prescribed emissions targets will not be exceeded. This is true even if regulators adjust future targets in response to early safety valve purchases, because there is no guarantee that prices will drop below the safety valve price in the future. Conversely, the emissions outcome could be guaranteed by implementing the reserve auction, but this mechanism cannot guarantee that prices will always remain below the political enforceability threshold.

Conclusions: Putting the Pieces Together

Ultimately, in order to arrive at a successful outcome, architects of climate policy will need to translate key strategic goals into specific cap-and-trade provisions. In this paper, we have highlighted two critical design objectives—price stability and cost viability—and shown how these objectives map onto various potential policy instruments. Attainment of the first objective (price stability) follows from the inclusion of temporal flexibility provisions, combined with limited regulatory authority to promote broader market stability through targeted adjustments to the offset supply. The second objective (cost viability), which is tantamount to long-term system durability and therefore a critical prerequisite to the first, relies on the inclusion of provisions that align the actual market outcome with the underlying political willingness-to-pay for climate change mitigation.

Given the dependence of all other policy objectives on the assumption of long-term policy durability, it makes sense to place this goal at the top of any list of key policy objectives. The simplest way to promote this outcome would be to impose a price ceiling at a value that would ensure alignment of the actual price outcome with the underlying political willingness-to-pay. Under a basic safety valve with a carefully chosen trigger price, enforceability would be preserved under all possible outcomes, and the expected emissions deviation from the prescribed abatement goal would be small (but not zero). If the implied tradeoff were ultimately decided in favor of greater quantity certainty, then other implementation mechanisms like the reserve auction might be considered instead. In this case, the total amount of abatement (rather than the price) would be guaranteed, and the chances of exceeding the ceiling price (as opposed to the emissions cap) would be low (but not zero).

Of course, long-term durability is a necessary but insufficient condition for price stability, which requires additional firm-level flexibility and a market dominated by forward-looking participants. Flexibility itself is relatively easy to implement through the construction of banking and borrowing provisions, assuming borrowing is given proper attention, as indicated above. When such flexibility falls short of what is required to keep the market within reasonable bounds, we have suggested that subtle regulatory adjustments to the offset supply could enhance price stability and increase the foresight with which firms make compliance decisions.

The basic list of design recommendations is therefore quite short: First, policymakers should allow firms to bank and borrow allowances in pursuit of the temporally efficient outcome, constraining borrowing with either a combination of quantity limits and interest or (preferably) with a single deposit/collateral mechanism. Second, policymakers should allow regulators to maintain limited discretionary authority over the offset supply to manage any remaining market dislocations, to enhance overall stability and to further support the conditions that increase the efficiency of firm-level decisions. Finally, policymakers should apply a price ceiling that is high enough to minimize the threat of interference with normal market operation but low enough to remain politically credible, taking into account the unavoidable tradeoff between price and environmental certainty in designing the specific implementation mechanism.

These three provisions—banking and borrowing, a sensible offsets discounting rule, and a well-chosen price threshold—would arguably constitute a complete set of policies to promote an efficient, stable and durable carbon allowance market. Of course, other combinations of instruments might be defensible as well, but architects of policy will need to evaluate any set of proposed mechanisms according to their abilities to achieve precise policy objectives. While the number of mechanisms utilized should be sufficient to meet the various aims of policy, any mechanism that is duplicative or fails to serve a clear purpose should be jettisoned. The bottom line is that a successful policy will be one that provides a clear mapping between its component parts and the underlying strategic objectives. In designing a system that must endure for decades, if not longer, thoughtful upfront engineering is one ingredient that cannot be discarded.

References

EIA 2009. Energy Market and Economic Impacts of H.R. 2454, the American Clean Energy and Security Act of 2009. Report SR/OIAF/2009-05. Washington, D.C.: US Energy Information Administration.

Fell, H., B. K. Mignone and A. Paul. 2009. "Designing Temporal Flexibility Mechanisms to Manage Speculative Risks," Working Paper.

Jacoby, H. D. and A. D. Ellerman. 2004. The safety valve and climate policy. *Energy Policy* 32: 481-491.

McKibbin, W. J. and P. J. Wilcoxen. 2008. "Building on Kyoto: Towards a Realistic Global Climate Agreement," Energy Security Initiative Policy Brief 08-01. Washington, D.C.: The Brookings Institution.

McKibbin W. J. and P. J. Wilcoxen. 2004. Estimates of the costs of Kyoto: Marrakesh versus the McKibbin-Wilcoxen blueprint. *Energy Policy* 32: 467-479.

Mignone, B. K. 2009. "Emissions Targets in Cap-and-Trade: Choosing Reduction Goals Compatible with Global Climate Stabilization." Energy Security Initiative Policy Brief 09-02. Washington, D.C.: The Brookings Institution.

Mignone, B. K. 2009b. "A Safety Deposit Mechanism to Contain the Economic Costs of US Climate Policy," Working Paper.

Mignone, B. K., M. D. Hurteau, Y. Chen and B. Sohngen. 2009. Carbon offsets, reversal risk and US climate policy. *Carbon Balance and Management* 4:3 (doi:10.1186/1750-0680-4-3).

Murray, B. C., R. G. Newell, and W. A. Pizer. 2008. "Balancing Cost and Emissions Certainty: An Allowance Reserve for Cap-and-Trade," Working Paper 08-03. Washington, D.C.: Nicholas Institute for Environmental Policy Solutions.

O'Neill, B. C. and M. Oppenheimer. 2004. Climate change impacts are sensitive to the concentration stabilization path. *Proceedings of the National Academy of Sciences* 101: 16411-16416.

Pirrong, C. 2009. "Market Oversight for Cap-and-Trade: Efficiently Regulating the Carbon Derivatives Market," Energy Security Initiative Policy Brief 09-04. Washington, D.C.: The Brookings Institution.

Pizer W. 2002. Combining price and quantity certainty controls to mitigate global climate change. *Journal of Public Economics* 85: 409-434.

Roberts, M. J. and M. Spence. 1976. Effluent charges and licenses under uncertainty. *Journal of Public Economics* 5: 193-208.

Stavins, R. N. 2000. "Market Based Environmental Policies," in *Public Policies for Environmental Protection* (P. R. Portney and R. N. Stavins, Eds.), Second Edition. Washington, D.C.: Resources for the Future Press.

Wigley, T. M. L., R. Richels and J. A. Edmonds. 1996. Economic and environmental choices in the stabilization of atmospheric CO_2 concentrations. *Nature* 379: 240-243.

BROOKINGS
The Brookings Institution
1775 Massachusetts Ave., NW
Washington, D.C. 20036
brookings.edu