



**Intertemporal Regulatory Tasks and Responsibilities
for Greenhouse Gas Reductions**

By Jeffrey A. Deason and Lee S. Friedman

**Goldman School of Public Policy
University of California at Berkeley**

**lfried@berkeley.edu
jadeason@berkeley.edu**

revised March 2010

This paper is forthcoming in The Journal of Policy Analysis and Management

Abstract

Jurisdictions are in the process of establishing regulatory systems to control greenhouse gas emissions. Short-term and sometimes long-term emissions reduction goals are established, as California does for 2020 and 2050, but little attention has yet been focused on annual emissions targets for the intervening years. We develop recommendations for how these annual targets—which we collectively term a “compliance pathway”—can be set, as well as what flexibility sources should have to adjust in light of cost uncertainties. Environmental effectiveness, efficiency, equity, adaptability, and encouraging global participation are appropriate criteria by which these intertemporal policy alternatives should be judged. Limited but useful knowledge about costs leads us to recommend a compliance pathway characterized by increasing incremental reductions along it. This can be approximated by discrete linear segments, which may fit better with global negotiations. While the above conclusion applies to any long-term GHG regulatory program, many jurisdictions will rely heavily on a cap-and-trade system and the same pathway recommendation applies to its time schedule of allowances. Furthermore, borrowing constraints in cap-and-trade systems can impose substantial unnecessary costs. To avoid most of these costs, we recommend that sources be allowed early use of limited percentages of allowances intended for future years. We also find that a three-year compliance period can have substantial benefit over a one-year period.

Keywords: climate change, cap-and-trade program, environmental regulation, emissions trading, greenhouse gases, carbon emissions markets

JEL Classifications: L51, L98, Q48, Q54, Q58, H23

Introduction

We consider a problem that confronts many jurisdictions that either have or are considering regulatory systems to reduce greenhouse gas (GHG) emissions. Many jurisdictions have passed or are considering passage of statutes that set targets for greenhouse gas reductions. These statutes typically specify reduction targets in specific future years (e.g., a target in 2020 and one in 2050) as a percentage of some past year's emissions. Such goal-setting stops well short of defining a limit on aggregate emissions over the full time period. We here consider the intertemporal tasks and responsibilities that attend to implementing such statutes. We focus upon "compliance pathways," or the time progression of yearly required emission reductions.

We use the state of California as an example. In its 2006 Global Warming Solutions Act (known by its legislative bill number AB32), California committed itself to reduce GHGs to the 1990 level by 2020. Its Governor has aggressively supported these efforts, and issued Executive Order S-3-05 that further commits the state to achieve by 2050 emissions equal to only 20% of the 1990 level. The California Air Resources Board (CARB) is the state agency with overall responsibility for developing the regulations to achieve these goals. It has adopted the value of 427 million metric tons of carbon dioxide equivalents (mmts CO₂e) as the official 1990 level of emissions.¹ That amount is therefore the 2020 goal and 85 mmts (20% of that number) is the 2050 goal. California's new regulatory system is currently under design, and will be implemented beginning in 2012 when we estimate emissions are likely to be around 537 mmts.²

This article focuses on *intertemporal* tasks and responsibilities, with the aim of recommending specific policy actions. These actions are by no means obvious for a number of reasons. First, it is not clear by what criteria proposed actions should be judged even if California is committed to achieving its goals no matter what the rest of the world is doing. Second, any single jurisdiction like California is a strategic player in a perilous game whose outcome affects the entire world and is determined by the worldwide and not local effort. Essentially, no one can win unless the whole world collectively reduces GHG emissions to prevent a temperature increase of more than 3 degrees centigrade by the turn of the century. This

¹ Carbon dioxide is responsible for the lion's share of anthropogenic global warming; however, there are many other greenhouse gases. Carbon dioxide equivalents convert the global warming induced by greenhouse gases in aggregate to the amount of CO₂ that would produce an equivalent amount of warming, thus putting all emissions on a common scale.

² In 2004 California emitted 480 mmts. If emissions grow at a slowed rate of 1.4% annually from 2004-2012, this leads to 537 mmts in 2012. Actual emissions growth from 1997-2004 was at 1.9% annually, but early actions under AB32 can be expected to slow this somewhat.

means emissions reductions for the world that are somewhat like those adopted by California.³ Third, we simply have no experience at managing this particular type of problem, and thus there are no pre-existing, satisficing bureaucratic routines under which this problem can comfortably fit. The next section describes features of the global warming problem that make it unique as a problem for regulatory management.

In our third section we examine features of the least-cost intertemporal path that will restrict aggregate emissions during a decades-long interval to meet a given environmental standard. No one knows the least-cost path in advance because many determinants of it will only be revealed in the future. This uncertainty motivates interest in market-based regulatory methods like cap-and-trade programs to determine many of the source-specific reductions. We consider regulatory policies that enable the actual emissions reduction path to come as close as possible to the least-cost one. One of these is the speed at which emissions are reduced. We find that under a fairly broad set of circumstances, the least-cost path will be characterized by increasing incremental reductions over time. This feature is thus a desirable characteristic of any long-term regulatory program to reduce GHG emissions, whether or not it is market-based.

Another policy focus is the role of saving and borrowing (i.e. the use of past and future allowances) in cap-and-trade programs. We find that allowing this type of intertemporal flexibility in the market is an important tool for achieving a least-cost mitigation path. These market-based adjustments are likely to be highly desirable in terms of overall cost management if they do not threaten the environmental integrity of the overall program. We find that there are reasons to limit borrowing, but not to the degree that characterizes existing and planned cap-and-trade programs. Because these limits constrain market adjustments to any time path of annual allowance issuances, they heighten the policy importance of making the issuance time path be one of increasing incremental reductions over time (holding constant the aggregate emissions permitted for the entire period). We consider several mechanisms that allow borrowing within periods of five years or fewer. We find substantial benefit to an advance auction each year of some future allowances with early use allowed—we use 20, 10, 5, and 5 % of the vintages for the next four years ahead as an example. We also find that there can be substantial benefit from having a three-year compliance period for trading allowances rather than one-year compliance.

Our fourth section addresses jurisdictions like California that do not constrain their implementing agencies to a fixed aggregate emissions level, but instead only provide them with a short-term and a long-

³ See, for example, Table 5.1 on p. 67 of IPCC (2007).

term target for specific years. This seems likely to be a common method for establishing emissions targets. In such a statutory structure, the compliance pathway serves dual roles. Not only does it potentially constrain the shape of the actual emissions profile (if saving and borrowing are restricted), it also determines the environmental standard itself by defining total emissions over the time period in question. Like many other jurisdictions, CARB will be relying upon a variety of regulatory programs with a cap-and-trade program as a significant component, and we consider the intertemporal choices open to CARB and review criteria for making these choices.⁴ We clarify that CARB has considerable discretion to choose compliance pathways that vary in terms of their environmental stringency, cost, and shape. We consider the emissions and cost differences among three realistic but quite different illustrative compliance pathways. We then consider briefly how other jurisdictions have specified compliance pathways and what intertemporal flexibility they have permitted their sources. For California we discuss how these choices compare to views of its “fair share” of emissions, as might be defined by global negotiations. We also consider how technological progress and global strategic considerations affect pathway choice.

In our fifth section we summarize our recommendations. We conclude that all of these considerations can be achieved by a plan to establish long-term goals in roughly ten-year increments, in which the goals for the next decade are set firmly and tentative goals for successive decades are announced but not finalized until approximately five years before the decade starts. We suggest a series of linear compliance path segments, each featuring a greater annual decline in the cap than the previous, as a simple mechanism that creates a pathway featuring increasing incremental reductions throughout the 2012-2050 time period. Our recommended cap-and-trade borrowing provisions should allow the market to adjust this pathway to keep it close to least-cost. We present this plan in the California context, but we believe that its general features are applicable to any long-term GHG emissions reduction policy.

Background for Intertemporal State Actions to Combat Global Warming

⁴ The Scoping Plan adopted by CARB in December 2008 specifies a variety of regulatory strategies to be used, including some command-and-control regulation that gives little discretion to emitters, programs restricted to specific sectors that give some flexibility to sources about how to achieve reductions, and a broad cap-and-trade program that by 2020 will include 85% of all California GHG emissions within it. CARB estimates that the regulations apart from cap-and-trade will achieve about 79% of the reductions required by 2020, with cap-and-trade achieving the other 21%. Should CARB have overestimated (or underestimated) what the regulations apart from cap-and-trade will achieve in the covered sectors, the cap will automatically ensure that market-chosen reductions make up the difference. For an analysis of the scope issue, see Friedman (2010).

Global warming is not a problem that falls geographically near the source of the pollution. Increased emissions from anywhere raise the threat of ecological harm all over the globe. It does not matter if the extra emissions originate in the U.S. or in South Korea; either way they have the same harmful and global effects. The overriding criterion for all policies aimed at global warming is effectiveness: limiting warming to a level at which the world's ecosystem is sustainable for the indefinite future. No single nation or state can do this on its own: even if the U.S. were to stop permanently all GHG emissions, business as usual (BAU) in the rest of the world would still cause unsustainable global warming.

Since its founding in 1988, thousands of scientists worldwide have contributed to the United Nation's Intergovernmental Panel on Climate Change (IPCC). Its 2007 report indicates that achieving a stabilization rate of 445-490 ppm of CO₂e by the next century requires annual worldwide emissions to peak by 2015 and decline by 2050 to only 50-85% of the 2000 emission levels. This would still imply an average global temperature increase of 2-2.4 degrees centigrade (IPCC, 2007). Even achieving this will likely lead to some irreversible impacts; the report cautions that 20-30% of species will be at an increased risk of extinction if global warming exceeds 1.5-2.5 degrees centigrade. This increases to 40-70% of species if global warming exceeds 3.5 degrees centigrade. Many experts think that emissions cannot be reduced rapidly enough to achieve a stabilization level less than 500 ppm, and suggest that a goal of 550 ppm by the next century may be the best that we can do.

Despite the critical need to reduce annual global GHG emissions by 2050 to something like 50-85% of the 2000 level, it is not at all clear that the world will do this. BAU paths have emissions increasing over time, not decreasing. While many countries of the world are working actively to reduce their emissions, important countries like the United States and China have (through 2009) refused to adopt specific reduction goals. To a country like China, rapidly improving its relatively low internal economic standard of living through economic development, appeals to halt its also rapidly growing GHG emissions seem unjust. It is hardly responsible for any of the CO₂e that is currently in the atmosphere; why shouldn't the developed countries that created the mess be the ones to clean it up? The failure of the United States to adopt specific reduction goals, knowing that it has been the major contributor to CO₂e, may seem less explicable. But these two extreme cases reveal an important aspect of the problem: determining just how much each country should reduce its future emissions to achieve the desired global result. There is

neither a benevolent dictator nor a world government that can impose a solution. Therefore we must find a solution that will be voluntarily adopted by virtually all. We refer to this as the “fair share” problem.⁵

Another complicating factor is the long time period before atmospheric GHG emissions dissipate. Carbon dioxide, the GHG responsible for most anthropogenic warming, takes 50-100 years to dissipate. Put differently, the carbon in the atmosphere now is roughly the sum of the carbon emissions over the past 75 years. Each year’s emissions, even if substantially reduced, still add to the existing GHG and increase global warming. Even if substantial worldwide absolute reductions were to begin now, our annual emissions will still be pushing the atmospheric concentrations to higher levels that would be better not to experience. In other words, it is critical that the world act decisively over the next 40 years and beyond to reverse global warming before its adverse effects can become unmanageable and irreversible.

Carbon dioxide is a “stock pollutant” because its damage is a function not of the emissions flow during a single year but of its accumulated stock in the atmosphere. In principle, one could try and estimate the marginal damage that will be caused by additions to the existing stock each year, and then restrict annual emissions to the level where the marginal cost of the last ton abated equals the marginal damage that ton would cause. There are excellent efforts to provide such estimates like Nordhaus (2008) and Stern (2007), but all recognize that there is great uncertainty especially about the marginal damage levels. In terms of actual policy decisions, estimates like these along with other scientific and political considerations lead policymakers to focus on quantitative emissions reduction targets to achieve over the next several decades. The sum of allowed emissions over this period is generally accepted by analysts as the best single measure of the stringency of the adopted targets, and then analysis can focus upon designing a regulatory system that will achieve this environmental goal at the least cost.⁶

In our work that follows, we use the sum of allowed emissions from 2012-2050 as a measure of overall environmental stringency. The next section derives some important results about the shape of the least-

⁵ At the time of this writing, the December 2009 Copenhagen conference just ended with no legally-binding emissions limits adopted by the world’s countries. China announced that it will reduce the carbon intensity of its emissions by 40-45% per unit of Gross Domestic Product from 2005 levels by 2020; since its GDP is growing rapidly its absolute GHG emissions will still grow. The United States is considering adopting a 2020 emissions goal about 5.5% lower than its 1990 level and a 2050 goal about 81% below the 1990 level.

⁶ See, for example, Leiby and Rubin (2001), Paltsev et al. (2007) and Stavins (2007). Paltsev et al. is evaluating a number of proposed U.S. cap-and-trade programs for GHG emissions, and writes (p. 3): “...a better measure of stringency [than the mid-century goal] is the sum of national emissions permitted between the start of the policy and mid-century. Stavins also writes (p. 16): “...the best measure of policy stringency may be the sum of national emissions permitted over some extended period.”

cost path for a given stringency goal, and the likely magnitudes of gains from allowing market-based intertemporal adjustments that achieve the same goal as a compliance pathway set by policy-makers. The following section considers that policymakers often incompletely specify the environmental goal, and provides guidance to the implementing regulatory authority about how to proceed in this situation.

Least-Cost Compliance Pathways

In this section we consider the intertemporal tasks and responsibilities of a GHG regulator assigned a specific stringency goal in the form of a total emissions budget over a set time period. We consider how a perfect cap-and-trade system would achieve a least-cost compliance pathway, in order to identify salient features of this pathway. In this model of a perfect cap-and-trade system, saving and borrowing are allowed unfettered without creating any difficulty. We then discuss why there are reasons to limit borrowing in actual cap-and-trade systems and we consider the cost implications of various saving and borrowing policies. We also consider how the borrowing limits bear on the regulator's choice of the time schedule for issuing allowances.

For a Least-Cost Pathway, the Present Value of an Additional Reduction Will Be Equal in All Years

We illustrate this feature of least-cost pathways in the context of a simple cap-and-trade program that allows saving and borrowing, and in which equilibrium allowance prices are the marginal reduction costs. Suppose the regulator establishes a yearly schedule of allowance issuances (that sum to the emissions budget) and we think of this schedule as an intended compliance pathway. The regulator may do its best to set that pathway in a way that minimizes costs, but it will likely fail to do so since it has only limited information about the cost of emissions reductions (the marginal cost of abatement curve). With both saving and borrowing permitted, however, the market will reallocate allowances over time to achieve a least-cost path.⁷

The market will reallocate allowances in order to make the present value of an allowance in any one year equal to the present value in any other year. Suppose, for example, the regulator committed to beginning with a very modest reduction in the first year (issuing a fairly generous number of allowances) followed by a substantial reduction in the following year (significantly fewer allowances). With no borrowing or lending, market participants would expect the (scarcer) second-year allowances to sell at a substantial

⁷ The benefits of allowing such flexibility can be substantial. For example, Ellerman, Schmalensee, Joskow, Montero, and Bailey (2000) estimate that intertemporal emissions trading reduced costs in the US Acid Rain Program by \$1.3 billion or about 7% of the total cost savings in its first 13 years.

price premium to those in the first year. Let's say that these expected allowance prices were \$30/ton for the first year and \$40/ton for the second. That means participants in year 2 expect that they will be better off to undertake reductions that cost, say, \$39 rather than have to buy a \$40 allowance.

If participants are allowed to save some allowances that they buy in the first year, this opens up a more cost-effective strategy: buy additional \$30 allowances in the first year and hold them to use in the second. This is much cheaper than paying \$39 for an actual reduction. But as many participants recognize this, the overall demand for first-year allowances increases (and correspondingly demand for second-year allowances decreases). With a fixed supply, the price of the first-year allowance must rise, and that of the second-year allowance falls. When will this stop? When it is no longer profitable for anyone to save additional first-year allowances. This will occur when both have the same present value:

$$P_1 = P_2/(1+r)$$

If, for example the first-year allowance price rose to and settled at \$34 with $r = .03$, then the second-year allowance price must be \$35.02.

Despite the regulator's decision to begin with only a modest reduction requirement, market participants will choose to reduce more in the first year in order to save allowances that can be used to supplement those available in the second year. The total emissions over both years will still be the sum of the allowances issued by the regulators (maintaining environmental integrity), but the actual compliance path will be different from what the regulators envisioned. This difference is desirable, as the market has acted to minimize the cost of compliance through saving.

In a system where saving and borrowing is not limited, and where borrowing does not create any noncompliance risk, the market can always reshape any regulator-specified pathway to achieve least cost. Unfortunately, the real world is more complex, and unlimited borrowing is problematic, as we will discuss. The resultant potential conflict between the regulator-specified path and the least-cost path that is environmentally equivalent is the crux of the more detailed examples that we develop below.

There are other important reasons why borrowing and saving occur, although they are not our focus. At any time, unanticipated events may arise that alter allowance demand in either (or both) the current and future periods. Some saving may arise in anticipation of these events, such as the "convenience yield" when it is less costly to supply from inventory than other methods (see Brennan, 1958). When

unanticipated events do arise, market participants recalculate the benefits to themselves of reducing emissions in the current period versus deferring reductions until the future, and in the aggregate additional saving or borrowing could arise in order to follow the (recalculated) least-cost path. For example, a substantial unexpected recession reduces economic activity and typically would thus also reduce allowance demand when it occurs.⁸ The unexpectedly low allowance price for this time would set in motion the same forces that we have just described. Participants would realize that they could take advantage of the temporarily low price to buy more allowances and save them for future use. This increased demand for saving would cause price to rise and partially offset the pure “recession effect.” While we do not include these effects in our estimates below, they may well be substantial.

The examples so far are of saving allowances, but it is easy to give examples where borrowing (if allowed) would save money. If the regulator requires very stringent reductions initially and then only modest additional reductions later on, market participants would borrow future (relatively inexpensive) allowances in order to use them instead of the high-priced current ones. The opposite of an unexpected recession is an unexpectedly strong, robust economy in which the demand is higher than expected for current goods and services and the emissions necessary to produce them. This causes current allowance prices to be unexpectedly high (relative to expected future allowance prices), and if allowed, participants will find it cheaper to emit more now by borrowing allowances in return for greater reductions later on when they are relatively inexpensive.⁹

One might think, incorrectly, that the least-cost compliance pathway would be to defer all reductions as far as possible into the future. This would be true if all reductions were equally costly aside from timing. However, because the marginal cost of reducing rises within any year with the quantity reduced, deferring too many reductions to the future would necessitate a steep climb up the marginal cost curve in future years. A simple example can illustrate. Suppose the goal were to reduce a total of 300 mmts over a 3-year period. Suppose also that 100 mmts could be reduced for \$30/ton, but the next 100 mmts would cost \$40/ton and the last 100 mmts \$50/ton. Table 1 shows 3 different compliance paths. Path 1 is linear, Path 2 defers all reductions to years 2 and 3, and path 3 defers all reductions until year 3.

[Table 1 about here]

⁸ This occurred in the EU ETS, in which allowance prices dropped from about €27 in 2008 to €13 in April 2009. The drop may not only be due to the recession, but perhaps to additional uncertainty about future EU reduction commitments.

⁹ Other researchers, for example Stavins (2007) and Murray, Newell and Pizer (2009), describe the ability of firms to respond more quickly than governments to the type of changes described here.

The linear Path 1 is much less costly than Path 2, and Path 2 is less costly than Path 3. This is because the linear Path 1 allows all reductions to be undertaken at \$30 per ton, whereas Path 2 requires some \$40 per ton reductions and Path 3 requires \$50 per ton reductions. Nor, in this example, would it be cheaper to front-load the reductions. A Path 4 (not shown) that had 150 mmt reductions in each of the first two years would require some \$40 per ton reductions and would have cost \$9.85 billion in present value terms.

The real market task is of course far more complex. The time frame is much longer, and marginal abatement costs are uncertain. Nevertheless, the market generally does the best that it can in trying to set a compliance pathway that equalizes the net present value of expected marginal abatement cost along it.

The Unknown Marginal Cost of Abatement Curve Shapes the Least-Cost Compliance Path

The last subsection illustrated why market participants will in general have incentive to use mechanisms of borrowing and saving in order to minimize the cost of complying with an emissions budget. In this subsection, we wish to clarify more generally how the shape of the marginal cost of abatement curve influences the shape of the least-cost compliance path. We assume here that an aggregate emissions budget has been set for a specified period.¹⁰

It will be helpful to introduce some notation. Denote $E(t_i)$ as the emissions in year t_i , with r as the discount rate. Let $MC(E(t_i), t_i)$ be the marginal abatement cost, with $\partial MC / \partial E < 0$ (i.e. it is less expensive to abate an additional ton when emission levels are high). The second term t_i by itself represents technological progress with $\partial MC / \partial t < 0$, meaning that over time the marginal cost of abatement at a given emissions level decreases (i.e. improved, lower-cost methods of abatement become available). **For any aggregate emission reduction to be achieved at the lowest present discounted value over time, the present value of the cost of the marginal abatement in year i must be equal to that of year j .**

$$(1) \quad MC(E(t_i), t_i)e^{-rt_i} = MC(E(t_j), t_j)e^{-rt_j} \text{ for all } t_i, t_j \text{ within the period}$$

One simple implication of this is that, **for any given aggregate reduction, annual emissions will decline over time along the least-cost pathway.** The discount factor will be a smaller number for $j > i$, and thus the marginal cost in year j must be greater than in i to make the equality hold. If there were no

¹⁰ In this context, earlier work by Cronshaw and Kruse (1996) provides a treatment of allowance banking, and Rubin (1996) extends this to include both banking and borrowing.

technological progress, the only way to make the equality hold would be to have lower emissions (higher abatement) as we go further into the future. Since technological progress tends to lower the marginal cost of abatement over time, this means that to make the equality hold future emissions must be even lower.

Another way to see the same point is to take the partial derivative of the present discounted value with respect to time and recognize that it is zero along a least-cost path:

$$\frac{\partial PDV}{\partial t} = \frac{\partial MC}{\partial t} e^{-rt} - rMC e^{-rt} = 0$$

This simplifies to a version of the Hotelling rule—**the proportionate increase in marginal cost from one period to the next will equal the interest rate:**

$$\frac{\frac{\partial MC}{\partial t}}{MC} = r$$

Expanding this to show the MC function:

$$\frac{\left[\left(\frac{\partial MC}{\partial E} \right) \left(\frac{\partial E}{\partial t} \right) + \frac{\partial MC}{\partial t} \right]}{MC} = r$$

The numerator must be positive for this equality to hold, but we know that $\partial MC/\partial E < 0$ and $\partial MC/\partial t < 0$. The only way the numerator can be positive is if $\partial E/\partial t < 0$, i.e. decreasing emissions over time. Solving the above equation for $\partial E/\partial t$:

$$\frac{\partial E}{\partial t} = \frac{rMC - \partial MC/\partial t}{\partial MC/\partial E}$$

Less immediately clear is whether the amount by which emissions decline along this least-cost path is increasing, constant, or decreasing over time. This corresponds mathematically to whether the partial derivative of the above expression with respect to t , or $\partial^2 E/\partial t^2$ is respectively negative (i.e. bigger decrements), zero, or positive (i.e. smaller decrements). This expression need not have the same sign over the entire time frame and range of emissions changes of interest to us, although for the most part we will simplify to assume that it does. It is difficult to make generalizations about the least-cost rate of emissions decline based only upon theory because it depends heavily upon both the shape of the annual marginal cost curve for GHG abatement and the rate of technological progress.

Nevertheless, we can at least clarify somewhat how the shape of the marginal cost of abatement curve influences the amount by which emissions should decline over time along a least-cost path. To increase the transparency of the next step, let us temporarily suppress the role of technological progress by holding

technology constant. That is, we consider the shape of the marginal cost curve using only existing technology. Then we calculate from above the second partial derivative (dropping out the technology term $\partial MC/\partial t$):

$$\frac{\partial^2 E}{\partial t^2} = \frac{[\partial MC/\partial E][r(\partial MC/\partial E)(\partial E/\partial t)] - [rMC][(\partial^2 MC/\partial E^2)(\partial E/\partial t)]}{(\partial MC/\partial E)^2}$$

Both long expressions in the numerator contain the same terms r and $\partial E/\partial t$ that can be factored out:

$$= \frac{r(\partial E/\partial t)[(\partial MC/\partial E)^2 - (MC)(\partial^2 MC/\partial E^2)]}{(\partial MC/\partial E)^2}$$

The signs of all of the terms in the above expression are known except for the last term in the numerator:

$$= \frac{(+)(-)[+ - (+)(\partial^2 MC/\partial E^2)]}{+}$$

This expression is unambiguously negative if $(\partial^2 MC/\partial E^2) \leq 0$. It can only be positive if $(\partial^2 MC/\partial E^2) > 0$ (corresponding to a rapidly rising marginal cost of abatement), and then it must be large enough so that the last term outweighs the first positive term in the brackets, rendering the sign of the bracketed terms negative.

What does this mean? Consider some common shapes of marginal cost curves: (1) constant marginal cost; (2) marginal cost that increases linearly; and (3) marginal cost that increases at an increasing rate. The constant marginal cost case has second derivative equal to zero, and thus induces a least-cost path featuring bigger emission decrements over time. However, in this case the solution to the least-cost path is determined by the boundary conditions (the emissions budget and the total time period) rather than equalizing present values along the path, as they cannot be equalized: **with constant marginal cost, it is always less expensive to abate in the future compared to the present. The least-cost path is simply to push all abatement as far into the future as possible, subject to staying within the aggregate emissions budget.** Starting from the initial year, undertake no abatement at all until the emissions budget for the entire period is reached. Then allow no further emissions until the end of the period.¹¹

Case (2), the linear marginal cost curve, also has second derivative equal to zero, and the least-cost solution is to have bigger emission decrements over time. There is more abatement in the early years compared to the constant marginal cost case, but the path might still be characterized as one that begins gradually with only small incremental abatement each year, and saves most of the heavy lifting for the

¹¹ Note that we do not consider constant marginal cost to be a reasonable empirical possibility in this application.

future periods. **If the marginal cost of abatement increases linearly, then the least-cost path will be one of successively greater incremental reductions along it.**¹²

Case (3) is the marginal cost of abatement curve that is increasing at an increasing rate: the cost of emissions reduction rises rapidly with greater reductions. This happens when the only ways to reduce emissions further become sharply more expensive, perhaps approaching the limits of our technical ability to reduce GHG in production and/or the willingness of consumers to forego any more of these goods and services. While this type of function is necessary for a reversal of the pattern of incremental reductions over time, it is not sufficient. Reversals are only caused by unusual shapes within this set (where the square of the first derivative is less than the marginal cost times the second derivative, as might be true at the elbow of a function that is fairly flat at lower abatement levels but then turns sharply upward). That is, **many steeply rising functions will still have the same pattern of increasing decrements of emissions over time. However, if the marginal cost of abatement is rising rapidly from a relatively high level but relatively low slope, it will cause the least-cost path in this range to be one of successively smaller incremental reductions.**¹³ While emissions levels will continually decrease over the years, as the level declines in this region so will the absolute size of each successive year's increment.

There is nothing inherently implausible about marginal cost of abatement functions that may feature small regions that would satisfy the conditions for decreasing incremental reductions. However, such functions are likely a small subset of the plausible MC functions. Moreover, even functions that feature this behavior seem likely to feature it only at specific points or regions along the curve. As such, we are comfortable concluding that, in general, least-cost compliance paths are highly likely to be characterized by increasing incremental reductions almost everywhere along their paths.

Intertemporal Flexibility Generally Has High Value

¹² There are cases between (1) and (2), in which marginal cost is increasing but less than linearly, that have this path characteristic as well. For example, the case with $\alpha = 1/2$ in the next note.

¹³ At this point, we have not identified any marginal cost functions that cause this reversal, and thus they may be unusual shapes. We considered, for example, the family of marginal cost functions for which $MC(E(t)) = [B - E(t)]^\alpha$ with B a constant baseline level of emissions and $\alpha > 0$. When $\alpha > 1$, this equation has the rapidly rising marginal abatement costs necessary but not sufficient for the reversal in pattern. For this function, $\partial^2 E / \partial t^2 = -(r^2 / \alpha^2)(B - E(t))$ which is always < 0 no matter how large α gets (and therefore the size of the emissions decrement along the least-cost path still gets larger over time). As another example, we considered the family of functions $MC(E(t)) = Ae^{bE(t)}$ with constants $A > 0$ and $b < 0$. This family has a linear least-cost path, with constant emission decrements ($\partial E / \partial t = r/b$).

We consider here the value of intertemporal flexibility for several possible shapes of the marginal cost of abatement curve and with various discount rates. These cases are illustrative, and generally indicate a substantial value to flexibility. For all of them, we use an emissions budget at a “California” level of 11,847 mmmts of CO₂e allowed from 2012-2050.¹⁴ We do not treat technological progress explicitly.¹⁵

We have examined a number of marginal cost of abatement functions in order to get a sense of the plausibility of the patterns that we observe. Figure 1 and Table 2 provide a representative summary. One set of results is for a linear marginal cost function that rises from \$0 to \$95 per metric ton over the range of emissions. Another set is for a log-linear function chosen to be more steeply rising (at an increasing rate) with marginal costs over the range from \$0 to \$181.¹⁶ A third set is for a step function that begins at \$30 per metric ton for the first 100 mmmts, and then increases by \$10 per metric ton for each additional 100 mmmts. This function does not rise above \$80 per metric ton because the maximum reduction cannot exceed our initial emissions level of 537 mmmts. Thus overall it is mildly rising like the linear cost function but in a deliberately discontinuous fashion. For each function, we calculated the least-cost path at both 3% and 7% real discount rates, and to calculate savings and borrowing, we compared each least-cost path with a benchmark linear regulatory path allowing the same overall emissions that has two segments in order to pass through the 2020 and 2050 California targets.¹⁷

[Figure 1 about here]

[Table 2 about here]

Consistent with our earlier theoretical analysis, the least-cost path for the linear marginal cost function at both 3% and 7% discount rates is characterized by saving in the early years and increasing incremental reductions as time progresses. These paths are shown in Figures 2a (3%) and 2b (7%). At the 3% rate, the

¹⁴ California has not actually set an emissions budget, but this level is the implied amount if proceeding by a path that has one linear segment from our estimated 2012 level to the 2020 target, and then a second linear segment from 2020 to the 2050 target. In any event, the value of the budget is immaterial to our results. For simplicity our calculations do not include saving or borrowing from pre-2050 to beyond it. A slightly more complex model could specify a terminal condition that assumes emissions beyond 2050 remain at the target level of 85 mmmts and compute any savings or borrowing that would carry over, but this would have negligible effects on our results and would not affect our policy conclusions.

¹⁵ We might expect that least-cost abatement will be pushed further into the future when technological progress is considered. However, the extent of this flattening will depend upon the nature of technological progress; if learning by doing is an important factor, incorporating technological progress might suggest more reductions up front. See Goulder and Mathai (2000).

¹⁶ We did not examine further very flat marginal cost functions, as it is transparent with these that the least-cost path is tilted heavily toward future-intensive reductions (and thus very substantial borrowing).

¹⁷ Draft California documents indicate that CARB is planning a linear compliance path over time, although this has not been formally adopted at the time of this writing. Similarly, proposed U.S. legislation such as Lieberman-Warner (S2191) and Waxman-Markey passed by the House (American Clean Energy and Security Act of 2009) features linear compliance paths.

initial reductions are quite substantial. There are cumulative savings that peak at 938 mmts in 2027 (8% of the total emissions budget), and then dissaving without ever entering a situation of net borrowing. The least-cost path is achieved at a savings of 7% over the regulatory schedule. At the higher 7% discount rate, there is also saving in the initial years but at a much-reduced level. The savings peak in 2016 at 138 mmts (1% of the emissions budget), then dissaving occurs with entry into net borrowing in 2023 that reaches a peak of 748 mmts (6% of the emissions budget) in 2040 and is then repaid by 2050. The least-cost path at a 7% discount rate is achieved at a savings of 6% over the cost of the regulatory schedule. In this case the higher discount rate does not particularly increase the value of intertemporal flexibility, although it shifts the pattern from one involving only saving to one involving fairly substantial borrowing.

[Figures 2a and 2b about here]

For the log-linear function chosen to illustrate a rapidly rising marginal cost of abatement, the least-cost path for both discount rates is also one of increasing incremental reductions over time. As with the linear case, the initial reductions are substantially greater than those along the regulatory schedule, so that there are substantial periods of savings. At the 3% discount rate, savings continue in each year until 2029 where cumulative savings peak at 1391 mmts (12% of the emissions budget). Thereafter there is dissaving until the emissions budget balances in 2050, so there are no periods of net borrowing. The savings of the least-cost path over the regulatory schedule are substantial: costs are reduced by 21%. At the 7% discount rate, there is still saving in the initial years but substantially less than with a 3% discount rate. Cumulative savings peak earlier in 2023 at 499 mmts (4% of the emissions budget), and then dissaving begins. In 2043, the cumulative savings are exhausted and a relatively brief period of net borrowing begins until 2048 when small savings come in to balance the emissions budget. Peak cumulative borrowing occurs in 2047 but only at 46 mmts (.3% of the emissions budget). Cost reductions are 7%.

For the step function, the least-cost path at both 3% and 7% is also one of increasing incremental reductions. However, there are no periods with positive cumulative savings. Because marginal costs are constant within each step, it is less expensive to defer “within-step” reductions until future years at the same step. At the 3% discount rate, cumulative borrowing reaches a maximum of 5% of the emissions budget (562 mmts) in 2033, and intertemporal flexibility reduces the present value of costs by only 2%. However, at the 7% discount rate there is much more borrowing, reaching 16% of the overall emissions budget (1936 mmts) in 2033. Intertemporal flexibility in this case reduces costs by a substantial 20%.

In sum, this subsection has shown that there is a high value to intertemporal flexibility measured relative to a regulator-set linear pathway for the issuance of allowances. While regulators may choose a different pathway (indeed, we shall recommend this shortly), our main point is that large uncertainties like the shape of the marginal cost curve make it impossible for anyone to set in advance a compliance path that will coincide with the least-cost intertemporal path. The large variation in the particular patterns of borrowing and saving shown in our calculations is caused by different possible shapes for the marginal cost of abatement curve and by different discount rates. The value of the intertemporal flexibility in our calculations ranges from two to 21 percent of the total cost of the benchmark regulatory-set linear path.

It is important to remember that the illustrative values shown here completely abstract away from an additional well-known source of value to intertemporal flexibility: macroeconomic fluctuations in the overall economy. For example, a recent paper by Fell and Morgenstern (2009) estimates the gains from intertemporal allowance trading in the U.S. economy due purely to macroeconomic fluctuations and finds that such trading reduces cost by about four percent. Thus we think the overall gains from intertemporal flexibility are at least in the six to 25 percent range.

All of the least-cost paths illustrated are characterized by increasing incremental reductions over time. The flatter the marginal cost curve, other things equal, the more valuable borrowing will be (and conversely for savings). This pattern is apparent across the borrowing (and saving) rows of Table 2, where the step function is flattest, then the linear function, and finally the log-linear function. There is less saving and more borrowing at the higher discount rate with each function (compare Figure 2a and 2b).

Achieving Gains in Light of Borrowing Restrictions: Three-Year Compliance Periods and Advance Use of Near-Term Future Allowances

Most existing and planned GHG cap-and-trade programs allow essentially unfettered allowance saving, but provide no or few opportunities for allowance borrowing. Saving has high value in many circumstances, but the existing and planned mechanisms are already designed to capture that value. The EU ETS system, for example, has been modified to allow saving across its phases but still effectively limits borrowing to one year ahead.¹⁸ In this subsection, we consider why the unfettered borrowing

¹⁸ In the EU ETS, allowances issued during the first compliance period (2005-2007) were not allowed to carry over and be used during the second compliance period (2008-2012). This caused a substantial problem, as allowance holders wanted to save the relatively inexpensive period 1 allowances to use in period 2 rather than purchasing period 2 allowances at the high prices expected for them. This regulatory ban on savings (across periods) caused the

permitted in the prior analysis is unlikely to characterize actual policies, and whether there are more limited but feasible borrowing mechanisms that could achieve most of the cost-reducing gains.

Aggregate allowance borrowing is a situation in which emissions in one year are greater than the number of allowances issued for that year plus older allowances still in savings, with the difference subtracted from future allowance issuances. California's Market Advisory Committee (2007, pp. 66-67) recommended against borrowing as possibly retarding environmental progress and technological advance, and to avoid the possibility that borrowed allowances will not be repaid. But the Committee also recommended that the compliance period for comparing an individual source's emissions with its submitted allowances be approximately three years.¹⁹ That is, a source would be required every three years to submit allowances to cover its emissions for the past three years. This is equivalent to allowing short-term borrowing within each three-year compliance period. The Committee recommended this to enable market participants to manage emissions levels in the face of unexpected short-term events (e.g. a year with unusually heavy electricity demand requiring more GHG allowances than expected).

When borrowing is limited, the market may be unable to achieve a given environmental goal (sum of reductions) at least cost. This heightens our interest in the specific initial, regulator-chosen time path for distributing the allowances. With no borrowing limits, the path's only significance is in defining the sum of allowances over the years; the market will reshape it to achieve least cost. But with borrowing limits, the market may be unable to adjust the path enough to achieve the overall environmental goal at least cost. The more front-loaded are the reductions on the compliance path chosen by regulators, the more likely that borrowing would be the market's method of adjusting to minimize costs, and the greater the likely divergence of actual compliance costs from the least cost way of achieving the same environmental goal. If regulators are uncertain about the optimal compliance path to set, they should avoid the front-loaded ones. Put more strongly, we have shown that the least-cost path is almost always characterized by increasing incremental reductions over time. With borrowing limits, it becomes important that the regulator-specified path also has this characteristic. For example, we are suggesting it will be unnecessarily costly to achieve any given emissions limit for the 2012-2050 period by issuing allowances in annual amounts that decline linearly over this period.

near complete collapse of allowance prices near the end of 2006 and throughout 2007, and retarded achieving actual emissions reductions. See King (2008), p. 70.

¹⁹ This has no bearing on the period for reporting emissions, which is expected to be no longer than annually for small sources and more frequently for large sources.

We have been referring to borrowing in the aggregate, defined as a reduction from scheduled future allowance issuances in order to cover a difference between current emissions and allowances submitted from the current year and from prior savings. To the extent that borrowing arguments focus on individual participants rather than just the aggregate, there may be no reason to set up a special allowance borrowing mechanism. Current allowance costs can be treated like the cost of any other inputs used by the source; they are simply one of the many costs of doing business. All the standard market borrowing mechanisms to assist a business in financing its operations exist; these should be sufficient to handle individual borrowing needs that arise due to allowance costs.²⁰

Therefore, we think the main reason for a system that enables allowance borrowing is for the concern that we are addressing: aside from transaction costs, aggregate borrowing can facilitate a less costly way to meet a given environmental goal than by strictly following the regulator-specified compliance path. No matter how wisely and carefully the regulatory compliance path is set, there is substantial uncertainty about the true shape of the marginal cost of abatement curve, the discount rate that sources will apply, and other circumstances that will inevitably change and cause some adjustment to the least-cost path. Knowledge about marginal costs is learned first through market experiences, and changing circumstances do not occur uniformly but affect different industries and different technologies at different times. It is the market rather than the regulator that is best suited to adapt to them.

Our illustrative calculations suggest that there are many plausible scenarios in which fairly substantial borrowing may characterize some portion of the least-cost path.²¹ Perhaps a good example that avoids extreme cases is the step function at 3% discount.²² Cumulative borrowing along the least-cost path reaches a peak of 562 mmts in 2033 (Table 2). While this represents only 5% of the total emissions budget, it is 201% of that year's cap of 279 mmts (the point on the benchmark regulatory-set linear path,

²⁰ One exception to this argument is if the transaction cost of obtaining current allowances becomes significant relative to obtaining a future allowance. There could be simultaneous saving by some sources and borrowing by others, and this may be less costly than if extensive search costs would be necessary for the borrowers to find and make matches with the savers. This is most likely to occur as the market for current allowances approaches the time for compliance (when most trades have already occurred), and thus would presumably be limited to some small percentage of that year's cap. It is perhaps a need like this that is addressed by the limited borrowing of the EU ETS system described by Trotignon and Ellerman (2008).

²¹ Details of calculations reported in this section are in Appendix 2 available on request from the authors.

²² We think this illustrative cost structure is a plausible range. It is roughly consistent, for example, with the estimates of McKinsey & Co. (2007) that the U.S. as a whole can achieve reductions of 40-60% from 2005 levels in 2030 at marginal abatement costs of \$50 with existing technology.

see Appendix 2). Peak borrowing for any single year along this least-cost path is 78 mmts, which occurs in 2026 when it is 22% of that year's cap of 359 mmts. Are there borrowing mechanisms that can help to achieve this type of intertemporal flexibility, without threatening other goals of the program?

We are respectful of a number of factors that make policy-makers reluctant to allow long-term borrowing. One is that we are in a world in which securing global cooperation to achieve worldwide reductions is paramount. To this end, it is of substantial symbolic importance that jurisdictions supporting global cooperation make real, verifiable reductions over time frames much shorter than the 39-year period we are discussing. Of the jurisdictions either already active in capping emissions or, like California, about to start, all either have adopted or are adopting short-term reduction goals, typically for 2020. That is, even if circumstances were like those in the high discount rate version of our step function, in which the least-cost path entails no actual reductions until 2026, we think this would not be tolerated. There would be too much skepticism that governmental commitment to real reductions is not credible.

A second factor is the uncertainty about global warming's physical effects, and the possibility that such effects could be catastrophic and become irreversible. This "fat tail" problem suggests much greater prudence than might be based on estimates of the most likely effects, and we think that restriction on long-term borrowing is an example of such prudence.²³ A third concern is one raised in the MAC report that unfettered long-run borrowing may reduce somewhat the incentive for the research and development necessary to support technological progress. While this is true of long-term borrowing for any economic activity, it is also true that ordinary market incentives for technological progress are generally too low because often the benefits of discoveries cannot be fully appropriated. If reasonable limits on long-term borrowing are paired with strengthened climate change R&D focused upon long-term solutions, that is perhaps a good tradeoff. A fourth possible concern, also raised by the MAC report, is that if borrowing in the aggregate is allowed from future entitlements, then there must be reasonable limits set on individual source borrowing in order to protect against market exit with unrepaid allowance debts.

The above considerations cause us to focus initially on mechanisms that generally limit borrowing to be within 5-year periods. One such mechanism is expanding the compliance period that sources have to

²³ In the presence of uncertainty and potential catastrophic outcomes of further warming, the "cost" of borrowing is potentially higher. If there is some unknown tipping point (to irreversible catastrophe), borrowing might cause us to go past it before we've had a chance to learn that it is there. The same emissions total with no (or less) borrowing gives the regulator more time to discover the tipping point and change its plans. See Weitzman (2009a and 2009b) for expositions of the fat tail problem applied to climate change.

submit or “true up” the necessary allowances (as discussed earlier). A second potential mechanism is releasing future allowance vintages to the market early and allowing early use of some portion of these vintages up to 4 years ahead to cover current emissions. One advantage of this is that it eliminates the concern about individual source borrowing and repayment: the individual sources cannot “owe” allowances as they must buy them. Thus there are no repayments due as a consequence of early use, only fewer allowances left for future use. We envision these advance sales as accomplished by auction.²⁴ Something like this needs to be done anyway, as it signals strong government commitment to the system, and as sources need guidance in the form of expected future allowance prices in order to make good investment decisions. Many may wish to buy allowances in advance as insurance against unexpected future price increases (e.g. if the economy heats up two-years ahead). However, our idea is not just that allowances be auctioned, but that sources be allowed to use vintages up to four years ahead to cover current-year emissions. Sources will know that any use of them for the current period means that there will be fewer of them to cover emissions a few years later.

In our models, the three-year compliance period provides valuable “borrowing” flexibility so that costs are significantly reduced from the one-year compliance period of our earlier examples. We illustrate with the step function at 3% discount with respect to the benchmark regulator-set linear path. Allowing unlimited saving and a 3-year compliance period saves \$2.131 billion of the \$3.567 billion excess cost of the scheduled path compared to the least-cost path, about 60% of the excess cost.

The additional mechanism of allowing some early use of future vintage allowances can further reduce these excess costs. Specifically, suppose that each year we auction allowances as follows: the balance of those for the current year, 20% of those one-year ahead, 10% of the two-year ahead, and 5% each of the three-year and four-year ahead (and perhaps some for further-ahead allowances as well). We calculate the effect of allowing this option and unlimited saving assuming the step cost function, 3% discount, and one-year compliance relative to the linear benchmark.²⁵ This leads to a present value of the cost of compliance

²⁴ There are other ways to accomplish the same thing. For example, suppose allowances are not sold by the government at auction but are freely distributed. Advance distribution of the same portion of future vintages with allowed early use would be equivalent. One could then rely on sales through third-party brokers in the allowance marketplace to get these to the sources valuing them most highly. Or the government could simply require that these early-distributed allowances be submitted for auction, with the proceeds going to the owners.

²⁵ To make this comparable with our other linear-path calculations that satisfy the 2012-2050 emissions budget of 11,847 mmts, we do not allow the 63.75 mmts of advance allowances for 2051-2054 to be used during the 2012-2050 period. Allowing their use would slightly increase the cost savings, and of course if this borrowing mechanism were actual they would be available for use.

equal to \$174.049 billion, or a cost savings of \$2.984 billion or 84% of the excess cost of \$3.567 billion. This compares to the 60% saving of the 3-year compliance period. The advance auction allows the least-cost path to be followed almost exactly up to 2024, at which time the limit on borrowing begins to cause somewhat larger reductions than in the least-cost case.

If we allow both limited borrowing from advance auctions and a 3-year compliance period, we can achieve some additional cost reductions. The present value of the reduction cost with both options is \$173.944 billion, a savings of \$3.089 billion or 87% of the excess cost. It is clear that the advance auction mechanism with limited borrowing is doing most of the cost-reducing work. Still, there is no reason not to do both. Figure 3 provides a graphic summary of this analysis. The contrast between the regulatory linear benchmark path and the unfettered least-cost market path is apparent, as well as the ability of our limited borrowing mechanisms to provide the flexibility to come closer to the least-cost path.

[Figure 3 about here]

This system also has great advantages as a “safety valve” that does not threaten the environmental integrity of the program.²⁶ Suppose there are circumstances that absent this mechanism would lead to a “price spike.” For example, perhaps an unusually adverse weather pattern arises that significantly reduces the supplies of a relatively clean biofuel that had been expected to be a major way of reducing GHG emissions this year. In the short run, there may be only expensive alternatives for reducing current year emissions. However, our mechanism could prevent most of the price spike by borrowing to increase the supply of allowances available to cover the current emissions. This allows time and flexibility to make the reductions in the next few years instead, when the biofuel supply may be restored or expanded and other less-expensive methods not available in the short run utilized. Our mechanism leads to a more gradual increase in the level of the price path, effectively spreading the risk from such events over the years. Furthermore, because early use of allowances necessarily raises somewhat the future prices, incentives for continued technological progress through research and development are if anything strengthened.

Of course there is nothing magical about the precise time frames for these borrowing mechanisms. A four-year compliance period for “truing up” allowances with emissions would allow greater cost savings,

²⁶ See Philibert (2008) for an assessment of “safety valves” and similar price caps and price floor ideas for climate policy. See also CBO (2009) for Congressional testimony on these by CBO Director Douglas Elmendorf.

to be weighed against a somewhat greater likelihood of difficulties enforcing compliance. It also might be possible to assign sources compliance periods that vary in length, in which the most dependable sources are allowed longer compliance periods.²⁷ Similarly, one could have higher or lower percentages of allowances auctioned in the advance sales, and could limit the borrowing to a greater or lesser number of years ahead. It is also good to keep in mind that our models, because they do not incorporate any uncertainty from macroeconomic fluctuations, are underestimating the true value of intertemporal flexibility and thus mechanisms that improve it.

Let us sum up. We have illustrated that, regardless of the time path for allowances specified by regulators, it is a near certainty that the least-cost path for emissions that achieves the same aggregate reduction will be substantially different. This is because of all of the uncertainties involved: the shape of the marginal cost of abatement curve, the pace of technological progress, and other changes in the economy that affect either or both benefits and costs of emissions reductions at any particular time. If the market is allowed to adjust by borrowing or saving allowances in the aggregate, it will substantially lower the present value of the cost of achieving the emissions reductions. All of the GHG cap-and-trade programs that are operating or are nearing operating allow savings with few if any restrictions. Borrowing, on the other hand, is generally not allowed or severely limited. We agree that unfettered borrowing would be problematic for a variety of reasons. However, we have shown that there is substantial cost-saving value to two mechanisms of borrowing that still limit it to be within a quite short time frame—roughly five years or less. One is each year to auction and allow the advance use of small portions of future allowances—we suggest something like 20, 10, 5, and 5 percent of the next four year’s vintages. The other is to have a multiple-year compliance period for truing up allowances—perhaps three-year compliance.

These borrowing mechanisms are important to include no matter what the shape of the allowance path specified by regulators. However, we also showed that, for quite a broad range of shapes for the marginal cost of abatement function, the least-cost path for a given emissions budget is characterized by increasing incremental emissions reductions. Moreover, for the reasons we described above, borrowing constraints are necessary and will likely limit market adjustments. If the regulator requires too little abatement now

²⁷ Our calculation of the three-year compliance period assumed that the periods are 2012-2014, 2015-2017, etc. up to 2048-2050. We then computed the least-cost way of meeting the scheduled reductions within the three-year period. This leads to a cost estimate for the illustrated step function at 3% discount of \$174.902 billion. However, all sources do not necessarily have to have the same compliance periods. For example, sources could be randomly assigned to one of the first three years to be the start of their future individual three-year compliance periods, so that the burden of compliance-checking procedures falls evenly over the years.

and too much in the future (relative to the least-cost path), the market can compensate by saving. However, if the regulator requires too much abatement now and too little in the future, borrowing constraints may prevent the market from fully compensating. Thus we also recommend that the issuance path for allowances set by regulators should feature increasing incremental reductions.

Jurisdictional Choice When Constrained Only to Meet Targets for Particular Years

To this point we have assumed that the total emissions budget has been fixed. However, the policy problem facing current greenhouse gas regulators is often different than this. Following the lead of the Intergovernmental Panel on Climate Change, many jurisdictions specify a short-term reduction goal and increasingly a long-term goal as well, and then proceed to develop regulatory programs intended to meet them. California's AB32 specifies that emissions in 2020 are to be reduced to the 1990 level of 427 mmts, and the Governor's Executive Order specifies that emissions by 2050 are to be no more than 85 mmts. Thus, California is not constrained to a sum of allowances; its constraint is to choose some pathway that passes through the targets. This leaves the actual stringency dependent upon the implementing regulatory body's chosen pathway shape. We have not found to date any analysis within CARB of how to choose the path shape that will connect its 2012 emissions level to the 2020 and 2050 targets, and we believe that its draft documents illustrate a tentative linear path because it is simple and transparent.²⁸ Thus we consider next not just our general shape recommendation to achieve a given degree of stringency, but additional factors that may help a regulatory body determine the level of stringency and the cost of achieving it.

We believe two criteria are particularly important for determining this choice: equity and efficiency. Equity in this case refers to whether the emissions budget implied by the pathway represents a jurisdiction's "fair share" of global efforts to reduce GHG emissions. Efficiency in this case means minimizing the cost of achieving the emissions budget. However, additional criteria matter as well. Environmental integrity is a criterion to ensure that actual emissions are in fact limited to be within the emissions budget; this has been of some concern, for example, in considering the extent to which offsets may be used as a means of compliance and whether those offsets are indeed additional and verifiable. Adaptability is another important criterion, particularly with regard to how changing knowledge may affect the appropriateness of long-term goals, and procedures for revising the regulatory system in light of

²⁸ For example, CARB issued on Nov. 24, 2009 a document entitled "Preliminary Draft Regulation for a California Cap-and-Trade Program." Subarticle 6 identifies how the schedule of annual allowance budgets will be set, and provides illustrative numbers that are a linear decline from 2012-2020, adjusted for the phasing in of different sectors into the cap-and trade program (e.g. transportation fuels become included as of 2015).

such new knowledge. Strategic considerations, in terms of engendering global cooperation, may also affect a jurisdiction’s willingness to commit itself to a particular path.

We first illustrate that a large number of different compliance pathways can meet the adopted targets of a jurisdiction like California. These pathways differ in both their emissions budgets and the cost of following them. To help choose among them, we next review existing cap-and-trade programs for lessons about compliance pathways. We then consider “fair share” views of environmental stringency, technological progress concerns, and the strategic context of inducing global participation. We make recommendations that apply to California as well as other jurisdictions with similar objectives.

Three Illustrative Paths

Our Figure 4 illustrates three different pathways that achieve California’s mandated goals. The blue line in Figure 4 is the linear pathway with two line segments that we have used in earlier sections as the benchmark. The incremental reduction from one year to the next is the same size along a segment (13.75 mmts per year from 2012-2020, and 11.4 mmts per year from 2020-2050).²⁹

However, nothing in the legislation or the executive order specifies that a linear pathway should be used. An alternative pathway shape is one that each year has a smaller incremental reduction than the year before. Illustrative of this is a pathway with constant percentage depreciation of the allowed emissions (CDE). Given our starting point, a reduction of approximately 2.8% per year will achieve the 2020 goal, and thereafter a reduction of 5.2% per year will achieve the 2050 goal. Such a pathway is shown as the two red line segments in Figure 4.³⁰ Note that compared to the linear pathway, CDE has greater initial reductions and then smaller reductions as one approaches the goal. Such a pathway results in lower total emissions than the linear path: from 2012-2050 the linear pathway averts 9096 mmts (or equivalently allows 11,847 mmts) but the CDE pathway averts 10,435 (allows 10,508), almost 15% more.³¹

²⁹ This pathway has the formula:

$$E_{t+1} = \begin{cases} E_t - 13.75 & t \leq 2020 \\ E_t - 11.4 & t > 2020 \end{cases}$$

³⁰ It has the equation:

$$E_{t+1} = \begin{cases} 537 \times (1 - .02825)^{(t-2012)} & t \leq 2020 \\ 427 \times (1 - .05238)^{(t-2020)} & t > 2020 \end{cases}$$

³¹ Appendix 2 tables contain the annual emission amounts E_t for each pathway. For each, we calculate the yearly reductions as equal to the 2012 level of 537mmts minus E_t . Thus a given total reduction also implies the “emissions budget” for the period (for our 39 years, total emissions in mmts equal $537 \times 39 - \text{total reductions}$). For other purposes, it can be important to calculate a growing BAU base, and define reductions as the BAU base minus E_t .

The compliance pathway might also be bowed in the opposite direction, to have increasing incremental reductions over time (seen earlier as the general shape for least-cost paths). An illustrative shape of this type is shown as the green line in Figure 4. This pathway has constant percentage appreciation of the reduction amounts (CAR) along its two segments. That is, the reductions grow by about 3% per year from 2012-2020, and they grow about 2% per year from 2020-2050.³² This pathway allows the greatest total emissions: only 8554 mmts are averted (and 12,389 mmts allowed). Still, it is as fully compliant with AB32 and the Executive Order as the other illustrative compliance pathways. Both the CAR and CDE paths are illustrative because each is a member of an infinite set of paths passing through the same targets that vary graphically only in the “arc” of the bows.

[Figure 4 about here]

The Present Value of the Cost of Compliance

If one criterion for choosing a pathway is environmental benefit, another one is cost. While no one yet knows how much it will cost to reduce emissions along these pathways, we illustrate the order of magnitude of the differences using the step function marginal cost structure introduced previously. We calculate the net present value of the cost of following each of the three illustrative pathways. For simplicity, we assume that the annual emissions limits are enforced with no intertemporal flexibility.³³

One measure of cost is the market value of the reductions. Within any year, the undiscounted market value will be the reduction for that year times a price that equals the marginal cost of abating the last ton.³⁴ Using a 3% real rate of discount, the least expensive pathway is CAR at \$215 billion. The linear

³²The formula for this pathway is:

$$E_{t+1} = \begin{cases} 937 - 400 * (1 + 0.03083)^{(t-2012)} & t \leq 2020 \\ 827 - 400 * (1 + 0.02081)^{(t-2020)} & t > 2020 \end{cases}$$

It was chosen from a family of equations that allow for the arc of the curve to be more or less bowed. The general equation for the family that will fit the second segment is:

$$E_{t+1} = 427 + s[1 - (1 + g)^{t-2020}]$$

In this equation, g is the yearly appreciation rate and s is a number that determines how bowed the arc is. The higher s is, the less bowed is the arc and the lower will be the g necessary to fit the bow to its two end points. We chose $s = 400$ to give realistic changes in emission reduction levels from year to year along the bow. One could generalize similarly with the CDE form, with $E_{t+1} = 427 - s[1 - (1 - d)^{t-2020}]$ for its second segment, although we use $s = 0$ because of precedent in the Regional Greenhouse Gas Initiative to be discussed shortly.

³³ Also for simplicity we assume that our cost schedule is exogenous to California’s choice of compliance pathway. However, as we discuss later, the compliance path itself can affect the degree and timing of technological progress that lowers the marginal cost schedule over time.

³⁴ Recall that in each year we define the reduction as the initial level of 537 mmts minus the compliance level. True reductions are greater, in that BAU would lead to growing emission levels each year (an increasing base to compare

pathway costs \$230 billion. The most expensive pathway is CDE at \$285 billion. Using a 7% real rate of discount, the numbers are lower but the relative rankings are the same (CAR \$89 billion, linear \$92 billion, and CDE \$116 billion).

An alternative measure is the social cost necessary to achieve compliance each year. That is, rather than valuing each year's reductions at the market price, we add up the cost of each metric ton reduced (the number of tons at \$30, the number at \$40, etc., discounted to present value). This gives a lower cost measure for each pathway, although the comparative results are similar. At 3% discount, the social cost of reductions along the CAR path is \$161 billion. Along the linear path this cost rises to \$177 billion, and it is \$213 billion along the CDE path. Again, using a 7% discount rate does not alter the relative rankings although the cost figures are lower (CAR \$69 billion, linear \$76 billion, and CDE \$92 billion).

To consider further how California might choose its pathway, it will be useful to review briefly the existing cap-and-trade programs to see what lessons they might offer about compliance pathways.

The Compliance Pathways of Existing Regulatory Programs

Several cap-and-trade regimes are in place or under development. Three U. S. markets for local and regional air pollutants have been operating for several years or more: the Acid Rain Trading Program, covering SO₂ emissions from fossil fuel-burning power plants in the 48 continental states; the RECLAIM program, covering NO_x and SO₂ emissions from power plants, refineries, and other industrial sources in the South Coast Air Quality Management District in California; and the seasonal (May-September only) NO_x State Implementation Plan (SIP) Call, covering (as of 2007) NO_x emissions from electric utilities and large industrial boilers in 21 northeastern and mid-Atlantic states plus the District of Columbia. Two markets for greenhouse gases are currently in operation: the European Union Emissions Trading System (EU ETS), covering large industrial emitters in the EU, and the Regional Greenhouse Gas Initiative

with allowed emissions). While this would tend to raise the marginal cost of compliance, this will be offset to an unknown degree by technological progress that lowers the cost to reduce emissions. An alternative way to think of our assumption of a constant marginal cost schedule is that it is equivalent to that of a growing base with technological progress that precisely offsets what would otherwise be the increase in marginal cost necessary to achieve the compliance level of emissions. Depending on the rate of actual technological progress, it could only partially offset the effects on marginal cost of an increasing base, or it could more than offset the increasing base effect (the latter has been the case for most exhaustible ores over long periods of time, see Nordhaus (1992)). Except in the case where each compliance pathway causes a different rate of technological progress, the marginal cost changes would be the same across all three compliance paths and thus would not affect the relative cost comparison. The market price in each period equals the marginal cost since it is the price that would result from a fixed level of emissions allowed each year (the compliance level) and competitive allocation of that number of allowances.

(RGGI), launched on January 1, 2009 and covering electric power generators in ten Northeastern states. A third GHG market, the Australian Carbon Pollution Reduction Scheme (CPRS) is proposed to commence July 1, 2010 pending legislative approval. It would be the most comprehensive to date, covering about 75% of Australian GHG emissions including those from the transportation sector. We discuss intertemporal aspects of each of these cap-and-trade regimes briefly.³⁵

The Acid Rain Trading Program was implemented in two phases, meaning that there was one step down in the cap that was planned from the start. In 2005 a federal regulation tightened the cap for 2010 and beyond, creating a second step down. Unlimited saving was allowed, including across program phases; even the statutory reduction is written in such a way that existing permits retain their value. Borrowing of allowances (the use of future allowances to cover today's emissions) was not permitted.

The RECLAIM Program featured a reduction in cap levels over time for both NO_x and SO_x, to a constant value beginning in 2003 (a target date for compliance with ozone requirements in the local Air Quality Management Plan). The compliance pathways were linear. The NO_x pathway had two linear segments: the cap decreased more quickly from 1994 to 2000 and then more slowly from 2000 to 2003. The SO_x pathway had only one linear segment. We have not uncovered any discussion of why a linear path was chosen. Significantly, RECLAIM did not allow saving or borrowing of emissions from one year to the next. During the first several years of the program, emissions were well below the cap and credit prices were low—an environment in which saving might seem attractive. In 2000, as the cap was beginning to constrain emissions more significantly, the California electricity crisis hit the market. As emitters had no saved allowances held in reserve, they found it impossible to comply with the cap, and credit prices in the NO_x market soared tenfold from 1999 to 2000.

The NO_x SIP Call program began in 1999 (then called the NO_x Budget Trading Program with 12 participating states) and was designed to impose a progressively more stringent cap over time. However, the actual cap levels have moved around substantially since 2003 as additional states have joined the program, and we have not determined the logic of the initial compliance pathway. The program allows saving, but places special restrictions on it to discourage a savings pool greater than 10 percent above the following season's emission target. If the savings pool exceeds this level, formulae go into effect that

³⁵ For good general reviews of cap-and-trade programs oriented toward their use for regulating GHG emissions, see Teitenberg (2003) and Ellerman, Joskow and Harrison (2003). Burtraw and Palmer (2008) have a very interesting analysis relevant to the issue of how allowances should be distributed in such systems.

effectively can require sources to provide two banked allowances rather than one to cover one ton of next season's emissions. No borrowing is allowed.

The EU ETS features distinct phases, the first of which has already drawn to a close. The first phase was a three-year "learning phase" from 2005-2007 with substantial noncompliance penalties of €40 per metric ton (about \$58). The initial caps for the learning phase were not very ambitious and allowance prices were low. Significantly, saving allowances within a phase was permitted, but not from one phase to the next. Borrowing is not allowed with one caveat: allowances for the next year are issued two months before allowances are due for the current year, and no restriction prevents the use of the next year's allowances to cover current year emissions. The compliance period for truing up allowances is one year.

Given that caps have tightened significantly in the second phase that covers 2008-2012, the prohibition on banking across the phases has resulted in prices of 2008 allowances that are much higher (around \$40 per ton) than the price of 2007 allowances (which went under \$1). Such price fluctuations are undesirable because they do not reflect any change in the social value of the emissions reduction they encourage. European carbon emissions did not suddenly become more damaging in 2008 when the prices rose, yet the incentive to avoid them abruptly changed. If EU member countries had been allowed to save their Phase I allowances for use in Phase II, they could have achieved the same environmental goal at a lower cost, and the prices would not have been so different between the phases. The restriction preventing saving across phases has been removed, so this problem will not affect future transitions across phases.

The EU ETS began as a relatively decentralized system where each member country produced a plan for its own cap level during Phases I and II. However, beginning with the eight-year Phase III period of 2013-2020, this will be replaced by a single cap on allowances for the whole EU. Furthermore, the EU has announced a linear reduction in the cap on allowances each year until 2020 to ensure "gradual and predictable reductions of emissions over time."³⁶ The size of the reduction was chosen to reach a 2020 target that is a 21.5% reduction of emissions from the 2005 level for the same covered sources, largely stationary sources that represent about 40% of total EU GHG emissions. The EU also offers a carrot to promote global-wide reduction commitments: it will increase its total 2020 GHG reductions (including sources outside of the cap-and-trade program) from its current level of 20% below 1990 levels to 30% if other developed countries commit to comparable reductions under a new global climate agreement.

³⁶ See Paragraph 13 of Directive 2009/29/EC of the European Parliament issued April 23, 2009.

The RGGI program offered from its outset in 2009 a clear compliance pathway design through 2018. The cap will remain flat at current emissions levels from program inception in 2009 through 2014. From 2015-2018, it will fall by constant depreciation of allowed emissions (CDE) of 2.5% per year in order to achieve a 10% reduction below the 2009 level. Unlimited saving will be allowed, and a three-year compliance period for turning in the appropriate number of annual allowances effectively permits some short-term borrowing. Furthermore, the length of the compliance period may be extended by up to three additional years if average allowance prices exceed \$10/ton (adjusted annually for inflation), effectively allowing additional borrowing in these scenarios. Despite the clear design, we have found no documentation of the rationale for the choice of this particular pathway shape. A comprehensive review will be conducted in 2012, to include consideration of whether additional reductions are warranted after 2018. The first quarterly auction for CO₂ allowances was held in September 2008, at which time 12.5 million allowances were sold at the price of \$3.07 per ton.

The proposed Australian CPRS cap-and-trade program is based on a long-term goal of achieving a 60% reduction from the 2000 level by 2050. The proposal also sets a 2020 goal of reductions that are 5-15% below the 2000 level, with 5% the minimum irrespective of the actions of other nations, and 15% if there is a global agreement in which all developed countries adopt comparable reductions. The government would specify annual caps 5 years ahead, and will also provide a further 10 years of guidance by specifying a “gateway” (a range) within which annual targets will be set. Each year the firm cap for the fifth year following would be announced, and the gateway updated every five years. Unlimited saving of allowances would be allowed, and sources will be allowed to borrow by the right to use the next year’s allowances for up to 5% of their current year allowance liabilities.

This summary of existing cap-and-trade systems highlights four points. First, saving plays a key role in averting large discrepancies in allowance prices from year to year, which otherwise can arise due either to program structure (EU ETS) or unforeseen events (RECLAIM). Second, borrowing is either not permitted or is sharply limited in all of these systems, with the main exception of RGGI’s three-year compliance period that can be made longer if triggered by high allowance prices. The no-borrowing restriction makes great sense for those pollutants that have short-term adverse health effects like SO₂ or NO_x, but this is not the case for carbon emissions. Third, none of these systems has a clearly articulated explanation for why its compliance pathway should be the preferred one. In fact, we have yet to uncover

any real evidence that the shape of the pathway was much more than an afterthought, and only Australia offers a long-term picture (beyond 2020). Fourth, Australia and the EU clearly have future global negotiations in mind.

What is California's "Fair Share"?

The literature on GHG reductions contains a large number of proposals for how to apportion reductions among countries or states. In this subsection, we consider whether California's reduction goals, if pursued along any of our three illustrative compliance paths, achieve its "fair share" of reductions. It should quickly become apparent that California's reduction goals implied by any of the three paths substantially exceed all but the most extreme concepts of fair share.

We consider the following proposed "fair share" standards:

Contraction and Convergence. Contraction is the process of reducing collective emissions to meet a concentration goal like 550 ppm. Convergence is the process of assigning the reductions to countries so that equal per-capita emissions are achieved by a chosen future year. This may be the most practical of the methods presented here, as it allows developing countries currently below the per capita target to have gradually rising emissions budgets.

Grandfathering. All countries retain their shares of current emissions, and reduce by a common percentage and time schedule to achieve a specific global reduction goal. This method would not be popular with developing countries, but it is closest to the methods used for free allowance distribution in existing cap-and-trade programs.

Global Preference Scores. Each country is assigned a share of a global emissions budget that is a weighted sum of two assignment methods: per-capita shares and grandfathered shares. The weights are determined by letting each country choose its preferred method, and defining the weights as the fraction of the world's population that has chosen each of the two methods. This method is perhaps best thought of as a democratic hybrid of the above two.

Our calculations for this subsection are in gigatons (Gt) of CO₂e, as is common in the literature (with 1 Gt = 1000 mmts). The 2012-2050 emissions on our 3 illustrative paths are: 12.4 Gt under CAR, 11.8 Gt under Linear, and 10.5 Gt under CDE. The fair share amounts for the same time period, based on a target atmospheric concentration of about 550 ppm, are: 15.8 Gt for Contraction and Convergence (with a convergence date of 2050), 28.4 Gt for Grandfathering, and 14.1-24.5 Gt for Global Preference Scores

(the lower figure if grandfathering is used to assign state shares within the U.S. and the upper figure if state shares are assigned on a per capita basis). Details of these calculations and the assumptions behind them are in Appendix 1. The important point of this exercise is that all three of our illustrative paths for California yield emission levels that are well under the “fair share” emissions levels assigned to them by these methods.³⁷ Thus we would not rule out any of the three paths by this criteria, and perhaps allow cost considerations to weigh more heavily than if these results were otherwise. If other jurisdictions adopt goals that achieve similar per capita emissions rates by 2050, they will likely meet “fair share” criteria as well regardless of the exact pathway shape chosen.

The Compliance Pathway and Technological Progress

There is much uncertainty about the relationships between climate change policy and rates of technological progress in addressing climate change. However, virtually all analysts believe that there is a relationship: the type of policies instituted affect the rate of technological progress, and generally increase it compared to no policies at all. Compared to unpriced and unregulated GHG emissions, charging a price for these emissions in the form of a cap-and-trade program is expected to induce technological progress. It does so for at least two reasons. One is that it provides incentives for private research and development activity intended to come up with cost-saving methods of reducing emissions. Very few analysts think this incentive provides a sufficient amount of R&D, however, because inventors cannot always capture the full benefits of their efforts. Additional government policies to increase R&D efforts, particularly those aimed at promising ideas not yet close to the commercialization stage, are often promoted by sponsoring basic research at universities and laboratories and through tax subsidies. The second reason is that the activity of reducing emissions itself causes “learning by doing,” so that some incremental improvements are routinely found and adopted over time (Goulder, 2004).

In order to understand the power of market incentives for technological progress, we shall turn away for a moment from the exhaustible atmosphere and consider the long-run record of exhaustible minerals provided through markets. Recall our earlier discussion of intertemporal allowance pricing: absent technological progress, an unfettered market is expected to cause the price of GHG allowances to rise over time with the interest rate. The same force applies to exhaustible minerals like aluminum or copper,

³⁷ If we used an extreme case of convergence in 2012 (rather than 2050) for Contraction and Convergence, meaning equal per capita shares from the outset, even then California’s fair share would be 11.0 Gts—just slightly below the Linear pathway—in a 550 ppm scenario. Of course, lowering the target atmospheric concentration would result in lower fair shares.

after accounting for any extraction cost: the rent or royalty portion of the mineral's price should rise with the interest rate.³⁸ If demand for these minerals is also increasing over time, that puts further upward pressure on price as the exhaustible supply becomes scarcer and scarcer. Yet economists have found just the opposite for most exhaustible resources characterized by increasing demand: over periods of 100 years or more, the real price of these commodities has fallen dramatically, often by 50% or more.³⁹ Why? Because the power of the market to find technical advances that economize on the use of the scarce resource has more than offset the upward price pressure.

If markets are created for GHG allowances, this does not mean that allowance prices will fall over the long run: we do not know, in this case, if technological progress will be strong enough to offset increasingly stringent limits on emissions. However, our interest is in the regulator's responsibility to create as effective a market as possible, so that it generates appropriate incentives for technological progress. It is not at all clear that existing cap-and-trade markets do this. We have already mentioned that the RECLAIM program does not allow saving or borrowing of allowances from one year to the next. This discourages investment efforts today that could have a greater return through saved allowances in the near term for use in the future (when allowances are expected to be more costly). As most proposed GHG programs, including California's, will allow saving, we do not expect this to be a problem.

A problem that characterized the early years of the EU ETS and discouraged technological progress is that there was very little knowledge of the compliance path beyond the phase ending in 2012. This was undoubtedly a function of the complex political negotiations between EU member countries necessary to reach agreement on a future path, but it added to the uncertainty about future allowance prices. This uncertainty discouraged R&D efforts that have a substantial part of their expected payoffs in these future periods. The EU's recent announcement of a specific path to 2020 is a great improvement, although some uncertainty remains due to its commitment to increase reductions significantly if a global agreement among developed countries can be achieved. We think it is very important that California and other jurisdictions provide a clear, long-run picture of the intended compliance path, although of course it needs to be adaptable to changing circumstances (e.g. new knowledge about the effects of global warming and the global efforts to prevent it). In other markets like those for exhaustible minerals, the relatively steady

³⁸ A mineral's price $P(t)$ at time t may be thought of as the sum of two components that can vary with time: the marginal cost of extraction $z(t)$ and the royalty $y(t)$: $P(t) = z(t) + y(t)$. See Chapter 19 of Friedman (2002) for an exposition.

³⁹ See, for example, Figure 5 of Krautkraemer (2005) or Figure 3 of Nordhaus (1992).

growth in market demand over the years gives a very important predictability to the market, even if any one year's demand is uncertain and subject to buffeting by the state of the economy or unexpected technological breakthroughs. But in the emissions markets, the number of allowances issued by the government determines the quantity demanded. The more predictable the government's actions, the more certain are the incentives for technological progress.

CARB knows with a high degree of certainty that it must meet the legislated AB32 target for 2020, and at a minimum should announce the allowed emissions for each year that will bring us from 2012-2020. CARB should always retain the ability to make adjustments due to unexpected emergencies, but there should be a very high degree of certainty that this path will be followed. We think CARB must go further than this, and we recommend a procedure along these lines: Announce a tentative interim compliance path for each successive decade that will bring us from 2020-2050. Adopt by 2015 a final compliance path for the 2020-2030 period, and confirm or adjust the tentative paths for the successive decades. Similarly, adopt by 2025 a final compliance path for the 2030-2040 period, and so on.⁴⁰ We recommend a similar approach for other jurisdictions, subject of course to the specific targets they may set.

While this kind of predictable long-run stability in the regulatory environment is extremely important, it could apply to any of our three basic shapes for compliance paths (CAR, linear and CDE). The main fact that we know about technological progress is that it lowers the cost of future emissions reductions relative to the present. If this effect derives from pure R&D, the least-cost path (for a given 2012-2050 cumulative reduction) would tilt somewhat away from the present and toward the future--in other words, toward the CAR shape. However, if it results in part from learning by doing, its effect on the pathway shape is unclear, as there is an incentive to undertake early reductions to gain the benefits of learning by doing. Thus the net effect of technological progress on the choice of pathway shape is ambiguous (Goulder and Mathai, 2000).

Strategic Consideration

It is critical that we do not overlook the most fundamental aspects of the climate change problem while we set the regulatory rules for particular jurisdictions like California. In particular, California's efforts (and those of the EU and other jurisdictions that are already taking actions) will be worth nothing unless

⁴⁰ The Warner-Lieberman bill S.2191 in the U.S. Senate specifies the exact allowance quantity for each year from 2012-2050, and proposes a Carbon Market Efficiency Board with the authority as a cost-relief measure to increase the number of allowances in any year through extended borrowing from future allowance allocations.

the world joins in and works for the global public good of mitigating climate change. Each jurisdictional GHG regulatory system must encourage and welcome other jurisdictions that adopt the necessary goals. This factor makes it difficult for individual jurisdictions to make credible commitments to long-term goals, as each jurisdiction's willingness to meet them depends on what other jurisdictions are doing.

We think that our pathway suggestion combines nicely with a strategic element suggested explicitly by the Australian model and the recent EU carrot (discussed earlier). The Australian model and EU carrot both promise a substantial increase in 2020 reductions (both by 10% from their baselines) if all other developed nations join in. We think this is the right idea, although perhaps too blunt. A related but more subtle idea is to approximate a long-run path that is intended to be CAR-shaped by a series of shorter-term linear approximations, each well within the "fair share" ranges. As more nations join in, the size of the incremental reductions along the next linear segment increases. This also allows adaptability to new knowledge as it becomes available. It is both economic and strategic.

Summary and Conclusions

This paper addresses the intertemporal responsibilities that GHG regulators like those in California face. We characterize the situation as one that begins with a short-term and a long-term goal and directs the regulators to set up a regulatory system that will meet them. In the California case, the regulatory system is to start in 2012, reducing emission to 427 mmts by 2020 and to 85 mmts by 2050 (equal to 1990 emissions and 20% of 1990 emissions respectively). There are an infinite number of compliance paths that can be taken to reach these goals, and the regulator must make significant choices about the intertemporal flexibility that sources will have. We apply multiple criteria in order to narrow the regulatory choices to a recommended subset.

The criteria we use include environmental effectiveness, intertemporal efficiency, equity in terms of a jurisdiction's fair share of global reduction responsibilities, strategic considerations relevant to inducing global cooperation, and adaptability to changing circumstances. Intertemporal efficiency raises a number of considerations, including the effect of the regulatory choices on technological progress, the value to sources of saving and borrowing, and the cost differences implied by alternative compliance paths.

One key responsibility of the regulators is to choose and announce a specific regulatory compliance path (the number of allowances to be issued annually from 2012-2050). We recommend that this path be set

with a high degree of certainty for the coming decade (subject only to unexpected emergencies), with preliminary paths announced for successive decades that are finalized approximately 5 years before each decade starts. The long-term path announcements take the place of trends in long-term aggregate demand in ordinary markets that are crucial to investment expectations and research and development efforts. GHG emissions markets will work less well than they should if characterized only by short-term emissions reduction goals. On the other hand, important new knowledge about global warming impacts and technologies available to reduce and mitigate them will become available in the future, and the regulatory system needs to be adaptable to this.

The environmental stringency of a compliance pathway for GHG emissions is determined by the sum of annual emissions. For a given degree of environmental stringency, we consider what is known about the shape of the least-cost compliance path that would achieve it. This depends on the unknown shape of the marginal cost of abatement curve. We consider a wide range of possible shapes for the marginal cost of abatement curve, and conclude that the least-cost compliance path is virtually always characterized by a shape that features increasing incremental reductions over time. Using discount rates of both 3% and 7% for the range of marginal abatement cost curves considered, we find cost savings of the least-cost path relative to a benchmark linear path that range from 2% of total costs to over 20% of total costs. These figures do not include the additional savings from flexibility to respond to changes in macroeconomic conditions, estimated by others as averaging about 4%. Therefore, the overall cost savings of moving from a linear to a least-cost path are likely in the 6-25% interval.

The shape of the least-cost compliance path would not matter to a regulator relying exclusively on cap-and-trade that allowed unfettered saving and borrowing. In this unrealistic case, no matter what the shape of the time path of allowances that the regulator schedules, the market will reset it to the least-cost path (with the same total emissions) by saving and borrowing as needed. But in reality actual greenhouse gas reductions are unlikely to precisely attain this least-cost path. For one thing, there are sources of GHG emissions that are impractical to include in cap-and-trade programs, and these sources may well be regulated in a fashion that does not allow their emissions the flexibility to reach a least-cost intertemporal path. Within a cap-and-trade program, we have seen there are good reasons for having some borrowing restrictions that constrain the market's schedule-resetting ability. However, borrowing restrictions may prevent the cap-and-trade market from attaining the least-cost path of increasing incremental reductions. *Thus for both market-based and other regulatory systems, cost considerations suggest that the time*

schedule for achieving aggregate reductions over a decades-long time frame should be one of increasing incremental reductions. However, actual practice and plans to date have not demonstrated awareness of this. California plans and the EU ETS are linear from 2012-2020, and silent so far about pathway plans beyond that. Proposed U.S. legislation features long term plans (to 2050) with linear reductions. RGGI features decreasing incremental reductions for its short-term plan (to 2018).

Within cap-and-trade programs for GHG reductions, the main reason for limiting borrowing relates to inducing the necessary worldwide participation. Leading jurisdictions, trying to demonstrate by example what all must do, must provide credible commitments to actual reductions. This means being able to demonstrate substantial real net reductions within a reasonably short time frame (e.g. 10 years). However, this does not imply that borrowing needs to be as severely limited as in current programs and plans.

This leads us to consider mechanisms that might provide valuable borrowing ability within five-year time frames. *We find that an advance auction and allowed early use of limited portions of future vintage allowances up to four-years ahead can capture a substantial share of the total cost-saving potential of unfettered borrowing (84% with our modestly-rising step function at 3% discount).* An additional way of effectively allowing some borrowing is to have a multi-year compliance or “true-up” period. California’s Market Advisory Committee recommended a three-year compliance period (i.e. a source would have to turn in the necessary allowances every three years). *Our calculations also suggest that there can be considerable value to the three-year compliance period (60% of possible cost savings with the same step function and 3% discount rate).* We know of no reason not to do both.

Our conclusions to this point are based on achieving a given level of environmental stringency, but the instructions given to regulators (for example, California’s instructions to CARB) do not specify this level of stringency. Rather, they specify target levels of emissions to achieve at specific points in time (in the case of California, for 2020 and for 2050). We illustrate three different paths that meet these targets but vary in terms of their overall stringency, cost and shape. The paths allow cumulative California emissions of 10.5-12.4 Gts of CO₂e during the period, and the ones with fewer emissions cost more to achieve. Another important consideration, from the perspective of inducing global cooperation as well as equity, is the amount of cumulative emissions that might be considered a jurisdiction’s fair share. We consider several different concepts of fair share that have received significant attention: Contraction and Convergence, Grandfathering, and Global Preference Scores. The “fair share” California emissions

budgets for 2012-2050 that we derive under reasonable application of these approaches range from 14.1-28.4 GtCO₂e. These are all well above the emissions levels of our illustrative paths. Thus we think California is likely to be regarded as exemplary if it follows any of the illustrative paths. Its 2050 target is aggressive, and California will have reduced its emissions considerably below the likely per-capita 2050 emissions worldwide by that time, in the process creating a non-carbon-intensive economy that will continue to benefit the climate into the future. Other jurisdictions that specify goals that attain similar per-capita emissions rates by 2050 will likely also be exemplary regardless of the precise pathway shape.

All of the above calculations assume a fixed technology for reducing emissions, but over long periods technological progress is one of the great drivers of all developed economies. A desired characteristic of any GHG regulatory system is that it generates R&D that will lower the cost of reducing emissions over time. While some of this can and should be publicly funded, it is also important that there be appropriate incentives for private R&D. We have emphasized the importance of specifying a long-run compliance path that is adaptable to new knowledge for creating a healthy private R&D and investment environment.

While there remains much uncertainty about how to best specify a compliance path, California's very strong reduction goals give it much flexibility—all of our paths that meet the targets generate emissions below the fair share amounts. On grounds of cost and intertemporal efficiency given borrowing constraints, we find that paths with increasing incremental reductions are most favorable. An approximation to such a path in decade-long linear segments may be the best choice overall, in that it preserves adaptability both to new knowledge and to what will hopefully be an increasing group of jurisdictions from around the globe that have agreed to fair share limits on their emissions. Thus we recommend that California regulators not only specify the compliance path from 2012-2020, but lay out a tentative long-run plan to 2050 with intent on specifying firmly by 2015 the path from 2021-2030. We also recommend that they plan to include some future vintages in annual auctions of allowances (whether or not any are freely distributed), and allow the early use of auctioned future vintages as much as four-years ahead. Finally, we recommend the three-year compliance period for truing up allowances. Other jurisdictions would similarly be well-served by a plan featuring these intertemporal characteristics.

Acknowledgments

We are grateful to the California Policy Research Center and the Toxic Substances Research Program, both of the University of California, for a grant that supported this research. We are also grateful for the comments and suggestions of participants at the IGERT Workshop, UC Berkeley, October 2008 and the International Conference on Environmental Policy in Seoul, South Korea, June 11-13, 2009, and the editors and anonymous referees of this journal.

References

- Burtraw, D., & Palmer, K. (2008). Compensation rules for climate policy in the electricity sector. *Journal of Policy Analysis and Management*, 27(4), 819-847.
- Brennan, M. J. (1958). The supply of storage. *The American Economic Review*, 48(1), 50-72.
- CBO (U.S. Congressional Budget Office) (2009, March 26). Flexibility in the timing of emission reductions under a cap-and-trade program. Testimony by Douglas Elmendorf before the Committee on Ways and Means, U. S. House of Representatives..
- Cronshaw, M. & Kruse, J.B. (1996). Regulated firms in pollution permit markets with banking. *Journal of Regulatory Economics*, 9(2), 179-189.
- Ellerman, A.D., Schmalensee, R., Joskow, P.L., Montero, J.P., & Bailey, E. (2000). *Markets for clean air: The U.S. Acid Rain Program*. Cambridge, UK: Cambridge University Press.
- Ellerman, A.D., Joskow, P.L., & Harrison, Jr., D. (2003). Emissions trading in the U.S.: Experience, lessons, and consideration for greenhouse gases. Report prepared for the Pew Center on Global Climate Change.
- Fell, H., & Morgenstern, R. (2009). Alternative approaches to cost containment in a cap-and-trade system. Resources for the Future Discussion Paper 09-14. Retrieved March 29, 2010 from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1392267.
- Friedman, L.S. (2002). *The microeconomics of public policy analysis*. Princeton University Press.
- Friedman, L.S. (2010). Should California include motor vehicle fuel emissions in a greenhouse gas cap-and-trade program? *Journal of Comparative Policy Analysis*. 12(3), 215-248.
- GCI (2005). GCI briefing: Contraction and convergence. Retrieved March 29, 2010, from <http://www.gci.org.uk/briefings/ICE.pdf>.
- Goulder, L.H. (2004). Induced technological change and climate policy. Report prepared for the Pew Center on Global Climate Change.
- Goulder, L.H. & Mathai, K. (2000). Optimal CO₂ abatement in the presence of induced technological change. *Journal of Environmental Economics and Management*. 39, 1-38.
- Intergovernmental Panel on Climate Change (2007). *Fourth Assessment Report*.
- King, M.R. (2008). Carbon markets and emissions trading: Lessons for Canada from Europe. *Policy Options*, 68-73.

- Krautkraemer, J.A. (2005). Economics of natural resource scarcity: The state of the debate. Resources for the Future Discussion Paper 05-10, Retrieved March 29, 2010, from <http://www.rff.org/rff/Documents/RFF-DP-05-14.pdf>.
- Leiby, P., & Rubin, J. (2001). Intertemporal permit trading for the control of greenhouse gas emissions. *Environmental and Resource Economics*, 19, 229–256.
- Market Advisory Committee (2007). Recommendations for designing a greenhouse gas cap-and-trade system for California. Retrieved March 29, 2010 from http://www.climatechange.ca.gov/market_advisory_committee/index.html.
- McKinsey & Company (2007). Reducing U.S. greenhouse gas emissions: How much at what cost?
- Murray, B., Newell, R., & Pizer, W. (2009). Balancing cost and emissions certainty: An allowance reserve for cap-and-trade. *Review of Environmental Economics and Policy*, 3(1), 84-103.
- Nordhaus, W.D. (1992). Lethal model 2: The limits to growth revisited. *Brookings Papers on Economic Activity*, 1992(2), 1-59.
- Nordhaus, W.D. (2008). A question of balance: Weighing the options on global warming policies. New Haven: Yale University Press.
- Philibert, C. (2008). Price caps and price floors in climate policy: A quantitative assessment. International Energy Agency Information Paper, OECD.
- Paltsev, S., Reilly, J.M., Jacoby, H.D., Gurgel, A.C., Metcalf, G.E., Sokolov, A.P., et. al. (2007). Assessment of U.S. cap-and-trade proposals. NBER Working Paper 13176. Retrieved March 29, 2010, from <http://www.nber.org/papers/w13176>.
- Rubin, J. (1996). A model of intertemporal emission trading, banking and borrowing. *Journal of Environmental Economics and Management*, 31(3), 269-286.
- Stavins, R.N. (2007). A U.S. cap-and-trade system to address global climate change. KSG Working Paper No. RWP07-052. Retrieved March 29, 2010, from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1026353.
- Stern, N. (2007). *The economics of climate change: The Stern Review*. Cambridge University Press.
- Teitenberg, T. (2003). The tradable-permits approach to protecting the commons: Lessons for climate change. *Oxford Review of Economic Policy*, 19(3), 400-419.
- Trotignon, R. & Ellerman, A.D. (2008). Compliance behavior in the EU-ETS: Cross border trading, banking and borrowing. Working Paper 08-12, MIT Center for Energy and Environmental Policy Research. Retrieved March 29, 2010 from <http://web.mit.edu/ceepr/www/publications/workingpapers/2008-012.pdf>.
- Weitzman, Martin (2009a). The extreme uncertainty of extreme climate change: An overview and some implications. Retrieved March 29, 2010, from <http://www.economics.harvard.edu/faculty/weitzman/files/ExtremeUncertaintyCliCh.pdf>.
- Weitzman, Martin (2009b). On Modeling and Interpreting the Economics of Catastrophic Climate Change. *Review of Economics and Statistics*, 91(1), 1-19.

Figures and Tables

Table 1: Choosing a Least-Cost Time Path

	Path 1		Path 2		Path 3	
	Reduction	Cost (\$ billions)	Reduction	Cost (\$ billions)	Reduction	Cost (\$ billions)
Year 1	100	3.00	0	0.00	0	0.00
Year 2	100	3.00	150	5.00	0	0.00
Year 3	100	3.00	150	5.00	300	12.00
Present Value		7.85		9.57		11.31

Figure 1

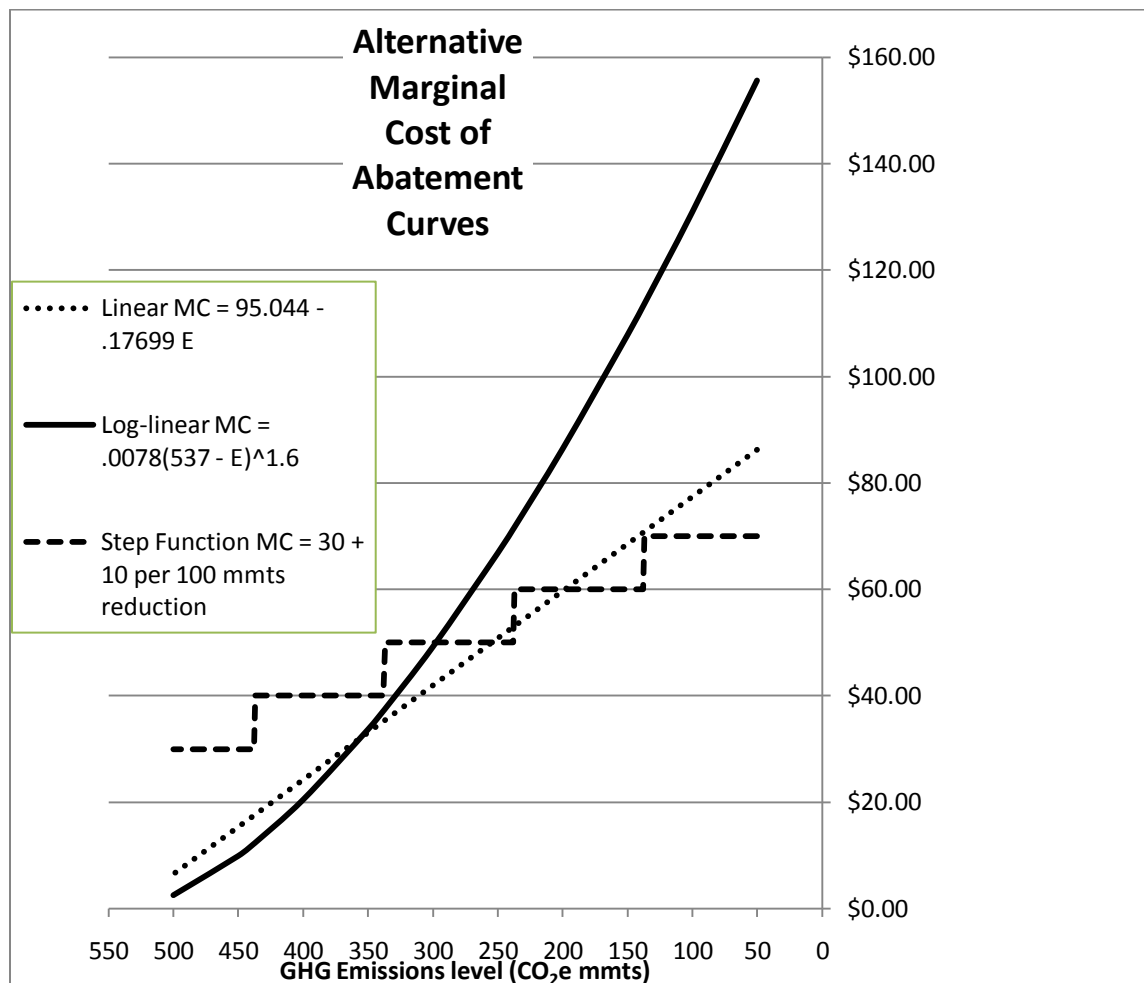


Table 2

The Least-Cost Path for Alternative Marginal Cost of Abatement Functions

(California Model: 2012-2050, 11,847 cumulative mmts of CO₂e allowed)

	Step Function	Linear	Log-Linear
Parametric Form	MC = 30 (+10 after each 100 mmts)	MC = 95.044 - .17699E	MC = .0078(537 - E) ^{1.6}
MC when E = 537	\$30.00	\$0.00	\$0.00
MC when E = 0	\$80.00	\$95.04	\$181.99
Least-Cost Path Features, 3% Discount			
Emissions in 2012	537	412	377
Emissions in 2050	0	147	211
Max. Cumulative Savings	0	938 (2028, 8%)	1383 (2028, 12%)
Max. Cumulative Borrowing	562 (2033, 5%)	0	0
Annual Reduction Increments	Increasing	Increasing	Increasing
Percent Cost Reduction	2%	7%	21%
Least-Cost Path Features, 7% Discount			
Emissions in 2012	537	489	447
Emissions in 2050	0	0	61
Max. Cumulative Savings	0	138 (2016, 1%)	499 (2023, 4%)
Max. Cumulative Borrowing	1936 (2033, 16%)	748 (2040, 6%)	46 (2047, .3%)
Annual Reduction Increments	Increasing	Increasing	Increasing
Percent Cost Reduction	20%	6%	7%

Figure 2a

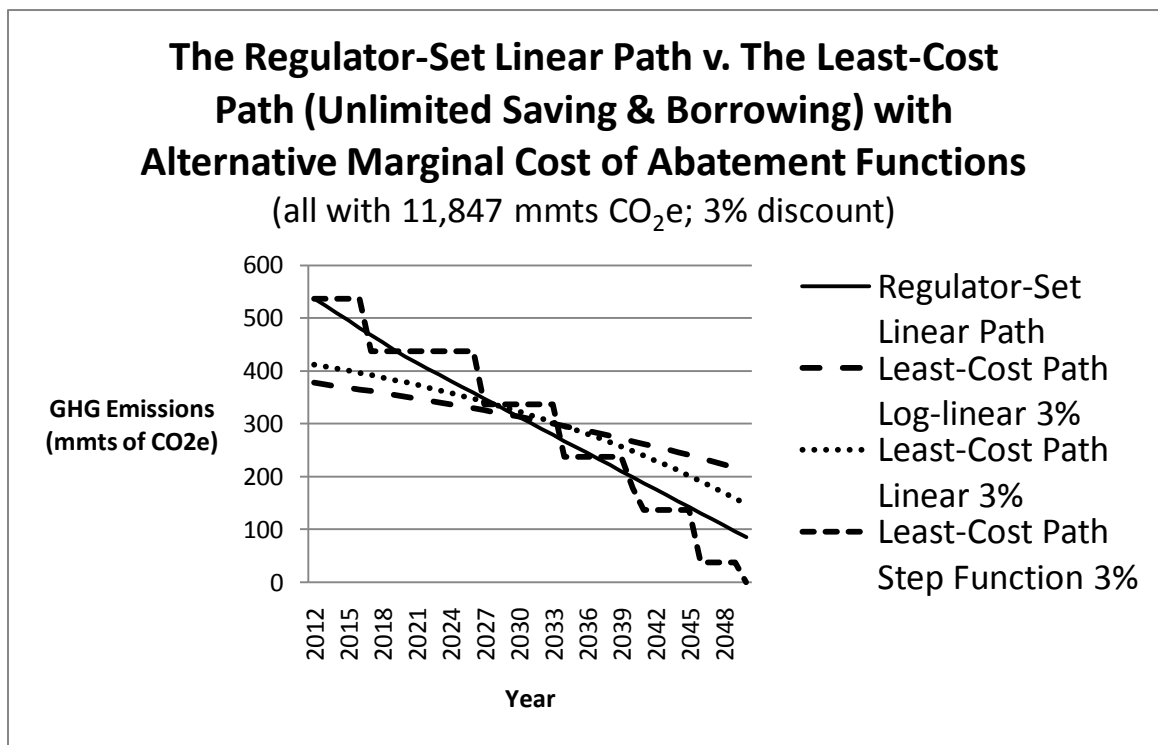


Figure 2b

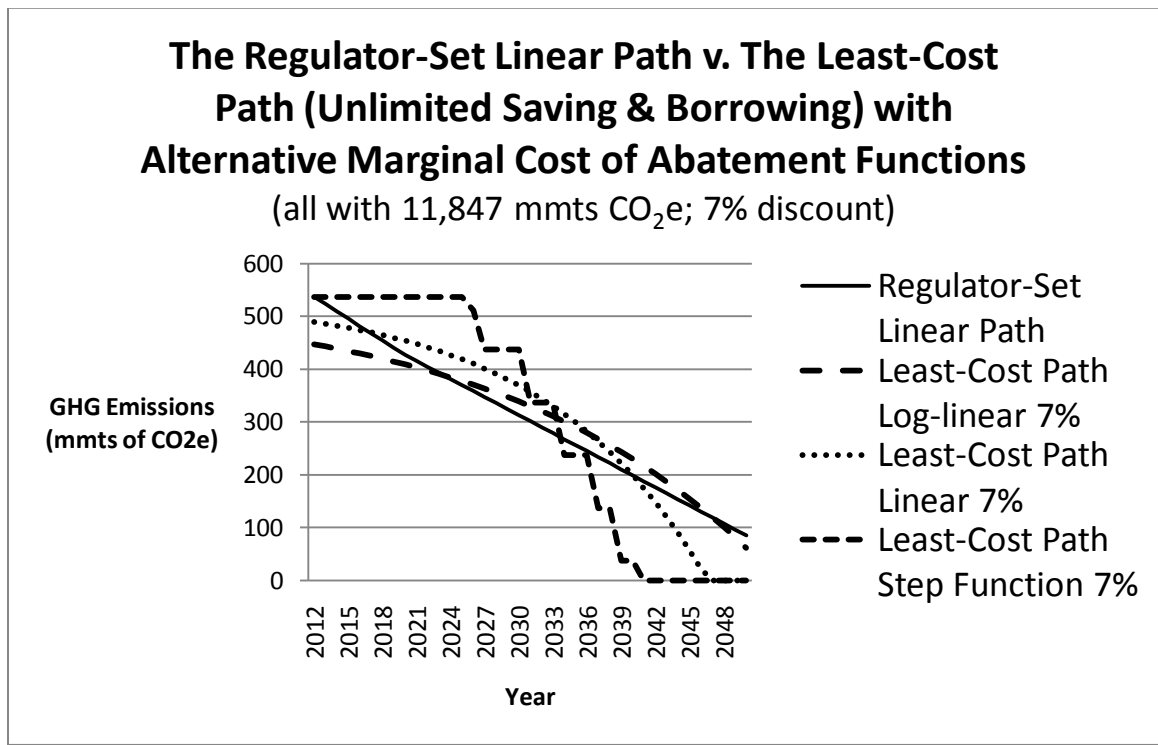


Figure 3

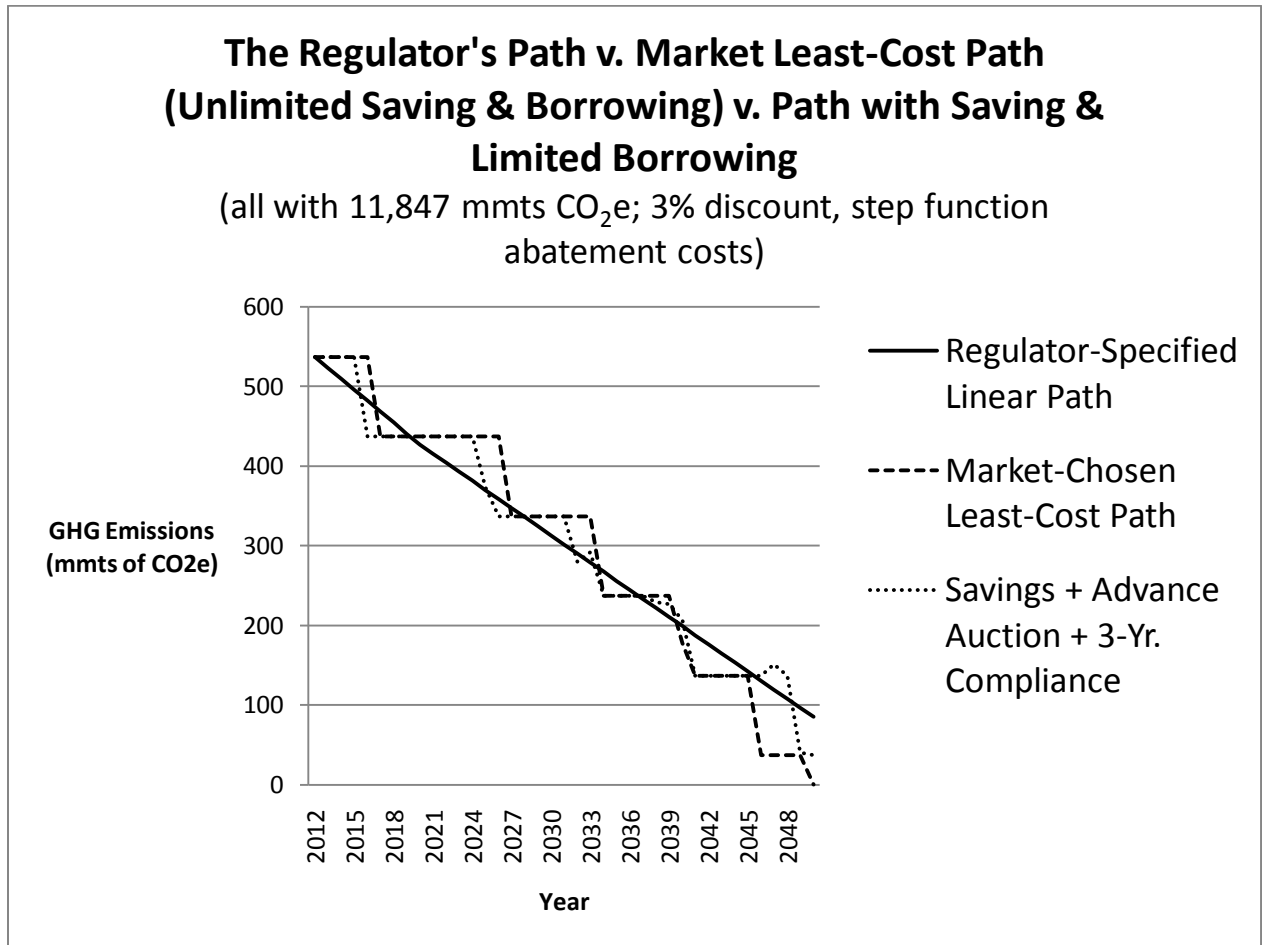
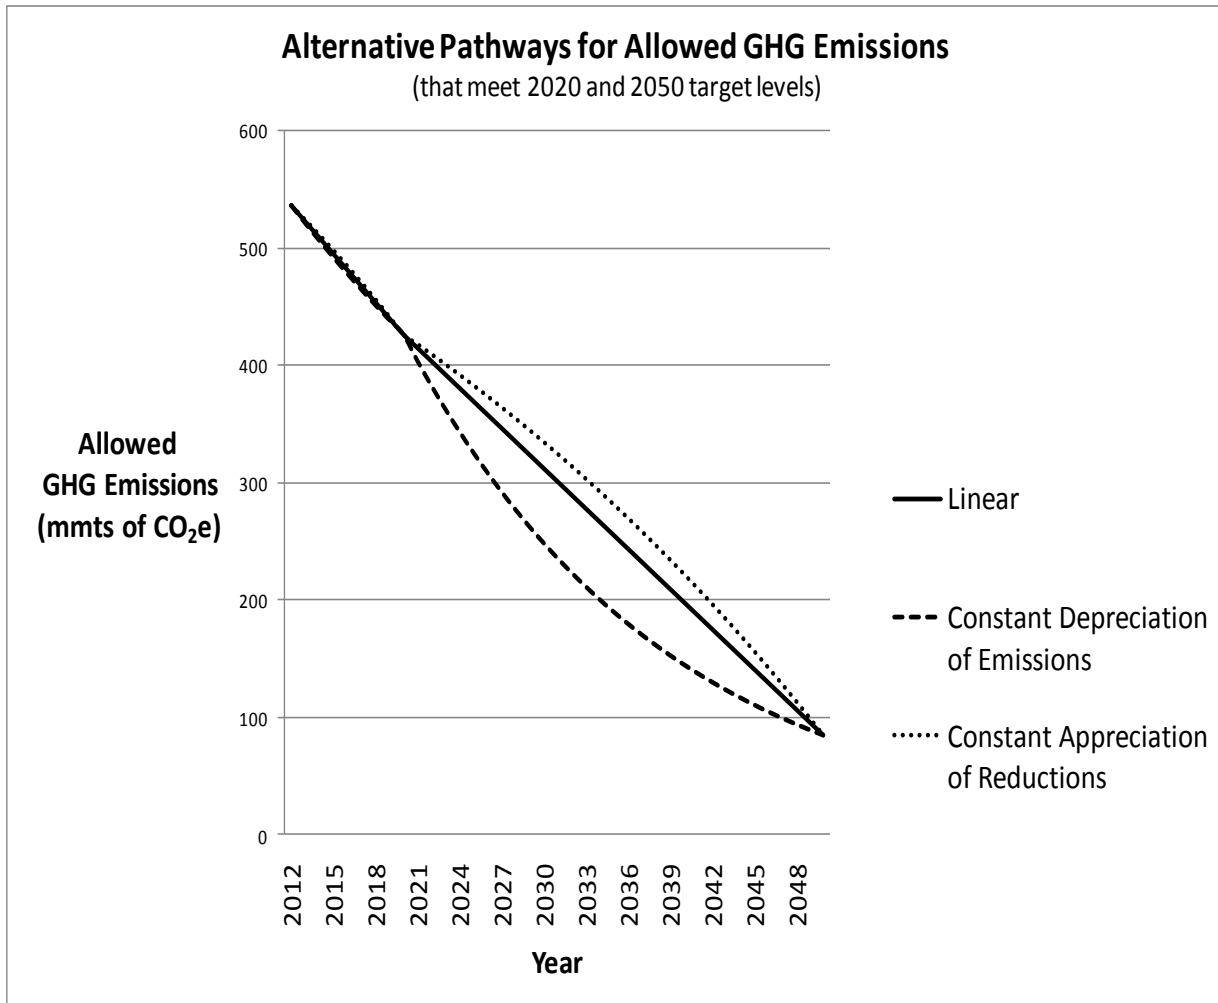


Figure 4



Appendix 1: Fair Share Calculations

A. Contraction and Convergence

Let us begin with a fair share concept that is one of the most appealing as a matter of practicality, simplicity and transparency: the “contraction and convergence” (C&C) standard described in GCI (2005). “Contraction” is the process of reducing collective emissions to meet a concentration goal, e.g., 550ppm. “Convergence” is the process of redistributing those emissions among countries to eventually attain equal per-capita emissions. Many of the fair share methods rely to some extent on the norm of equal per-capita shares, particularly in the long-run when unequal economic circumstances across jurisdictions might be reduced and all countries developed.

Under the basic approach, one must pick a year by which all countries will agree to converge on an equal per-capita allocation. One must also choose a base (starting) year, and a global emissions target in the year of convergence. All countries’ emissions allowances then change linearly over time from their current levels to an equal per-capita share of the chosen target in the year of convergence. While for most countries this will be a gradual decline in emissions, a number of less developed countries with current per capita emissions below the target for convergence would be allowed an emissions budget that gradually rises. Some authors, including GCI, recommend using the population shares as of the base year to avoid encouraging population growth; others recommend against that, reasoning that this incentive is small compared with all the other incentives involved and that such action effectively punishes future inhabitants of countries whose current growth rates are high.

For purposes of our exercise, we set the base year at 2012 (when California’s new regulatory system goes into effect) and the year of convergence at 2050 (the target year for attaining California’s long-run goal of emissions at 20% of the 1990 level). We specify the global emissions target in 2050 at 53.11 GtCO₂e, a level calculated from IPCC data on Scenario B1, which involves stabilization near 550 ppm.⁴¹ We used 2000 population levels: 6.122 billion for the world (from the same IPCC Scenario B1 Image) and 33.87 million for California from the U.S. Census, yielding a population share just under 0.6% for California.

⁴¹ We used the IPCC Scenario B1 Image figures for the GHGs included in the calculation of California’s emission goals (CO₂, CH₄, N₂O, CFCs, PFCs, and SF₆), converting by using the global warming potentials from the IPCC Second Assessment Report (as specified in the Kyoto Protocol) into CO₂e. These figures are commonly expressed in metric gigatons (Gt) and megatons (Mt), where 1 gigaton equals 1000 megatons and 1 megaton is one million metric tons (1Mt = 1mmt).

Under these assumptions, California's emissions allowance in 2050 would be 293.8 MtCO_{2e}. However, the Governor's 2050 target is 20% of 1990's 426.6 MtCO_{2e}, for a considerably lower value of 85.3 MtCO_{2e}. In other words, California's actual 2050 target is a very ambitious target compared to this standard. The C&C standard would only require California to cut back to 69% of its 1990 level, whereas the adopted goal is 20% of the 1990 level. The C&C goal may seem modest only because California starts from a relatively "clean" position for a developed jurisdiction. In 2020, allowable emissions under C&C with the above assumptions lie along a line segment between 2012's emissions (estimated at 537 MtCO_{2e}) and the 2050 value, for a 2020 target of 441.9 MtCO_{2e}. The AB32 target of 426.6 MtCO_{2e} is again lower, albeit only slightly.

Under the assumptions above, C&C yields a total CA emissions allowance of 15.8 GtCO_{2e} from 2012-2050. This is substantially above the 10.5-12.4 GtCO_{2e} range from our three illustrative compliance paths. Put differently, any of the three shapes would more than satisfy environmental equity if California's responsibility from a global perspective were judged by the C&C standard with 2050 as the date of convergence.

Of course there is nothing magical about 2050; one can imagine the date of convergence being either earlier or later and that would change the fair share calculations. We are reluctant to consider later dates because we think it critical that emissions globally are under control by this time. We do consider earlier dates, recognizing that they may be impractical (in terms of actual global accomplishment) unless new evidence propels the world to act more quickly than suggested by the most recent IPCC assessment. Moving the date of convergence forward to 2030 tightens California's emissions allowance to 13.4 GtCO_{2e}, still higher than any of our three illustrative paths. In the extreme case of convergence in 2012 or earlier (effectively meaning a uniform per capita standard around the globe from the outset of the program), California's cumulative emissions allowed from 2012-2050 would be 11.0 Gts, slightly lower than the allowance produced by a linear pathway.⁴² We think this last calculation simply shows the aggressiveness of the California standards, as we think the likelihood of global political consensus and action around this norm is quite low.⁴³

⁴² We also varied the population assumption, and rather than using the fixed base we used projected population over time. This increased California's allowed emissions, as its population is expected to grow faster than that of the rest of the world.

⁴³ Another way to see the stringency of this particular concept is to calculate the fair share amount for the U. S. as a whole and compare it to pending U.S. legislative proposals. The U.S. fair share amount by this definition is 91.39 GtCO_{2e} from 2012-2050. The most stringent U.S. proposal to date is Lieberman-Warner (S2191) that permits

B. Grandfathering

A strict grandfathering approach would require all countries to reduce current emissions by an equal percentage to attain some specified emissions schedule. All countries therefore retain their share of current emissions. Here we must choose the emissions schedule and the date at which “current emissions” are evaluated.

To be consistent with the approach chosen at Kyoto, suppose that the base year is 1990. In the year 1990, global emissions were 36.4 GtCO_{2e}. CA emissions were 426.6 MtCO_{2e}. CA’s share of global emissions is thus 1.17%.

Using IPCC Scenario B1, we estimate that total global emissions from 2012-2050 will be 2426 GtCO_{2e}. Applying CA’s 1.17% share yields a CA allowance of 28.4 GtCO_{2e}, more than double the allowance along any of our three illustrative paths. Indeed, since 2050 emissions are higher than 2012 emissions in Scenario B1, a simple freeze of emissions at the 2012 level would be allowable, so any pathway that reduces emissions would more than meet this criterion.

We have included this calculation because it is most like the methods that have been used to decide upon allowance distributions in the existing cap-and-trade programs. Despite these precedents, we note that no serious support seems to exist for a pure grandfathering approach at the global level (in which shares must be assigned to jurisdictions in very different circumstances from one another, unlike the units in the existing programs).

C. Global Preference Scores

The global preference scores method (GPS) is a weighted sum of two methods: per-capita and grandfathering. To derive the weightings, each country chooses which method it prefers. Then these preferences are aggregated based on the population of each country. The weightings for each method thus correspond to what fraction of the world would (indirectly) choose that method. GPS results in approximately an 8.4% U.S. share. 8.4% of the emissions in Scenario B1 from 2012-2050 is 203.8 GtCO_{2e}.

131.67 GtCO_{2e} during the same period. Reilly et al (2007) calculate the sum of allowances in the same period for 7 other U.S. legislative proposals. The range of these is from 148-306 GtCO_{2e}; see Table 3, p. 13.

We can use 2000 populations to assign California a per-capita share. Using census data, the U.S. population in 2000 was 281.4 million. As noted above, the California population in 2000 was 33.87 million. CA would thus have a share of 12.0% of the U.S. allowance, or 1.01% of the world allowance. This is comparable to but slightly smaller than the 1.17% share derived using the grandfathering approach. The resultant CA share from 2012-2050 is 24.5 GtCO₂e, again achievable by any conceivable non-increasing emissions pathway.

Of course, perhaps we should not assign U.S. shares on a per-capita basis. A grandfathering basis within the U.S. itself would be less favorable to California, as California emissions per capita are currently low relative to the U.S. as a whole. While we doubt that California's earlier efforts to stay clean would not be factored in to its U.S. share calculation, let us ignore this and consider a strict grandfathering approach using the usual 1990 as the base year. In 1990 CA's emissions were 426.6 MtCO₂e. U.S. emissions were 6.148 GtCO₂e. On this basis, CA would have 6.94% of U.S. emissions, or 0.58% of world emissions, for a share of 14.1 GtCO₂e. This is slightly under the amount that would be allowed by the C&C method, but still substantially more than enough to follow any of our three illustrative paths.

There are of course many other possible ways to assign fair shares, including those that would make use of data on economic circumstances like GDP, and those that assess historical responsibility.⁴⁴ Moreover, there are different concepts of the total reductions that should be apportioned on some "fair share" basis. Were we to take the recommended reductions of the Stern Review, for example, the suggested reductions would be much greater and California's reduction goals would not look as ambitious. On the other hand, if we used a different IPCC scenario, the suggested reductions would be less. Our purpose is not to resolve these complex issues, but merely to give some perspective to the goals that California has adopted. We think it is clear that California's adopted goals must be viewed as an extraordinary effort, regardless of which of the three path shapes it chooses for its compliance path. This is especially true given the target set for 2050, which is significantly lower than any of the fair share paths require. A California economy that has been reconfigured to achieve this emissions target in 2050 will be very well positioned to continue to serve as an exemplar in the years ahead by emitting less than its fair share.

⁴⁴ A good summary of many of these is contained in the IPCC Fourth Assessment Report, Chapter 13, Table 13-2.