HARVARD

JOHN M. OLIN CENTER FOR LAW, ECONOMICS, AND BUSINESS

TAXES, PERMITS AND
CLIMATE CHANGE

Louis Kaplow

Discussion Paper No. 675

08/2010

Harvard Law School
Cambridge, MA 02138

Taxes, Permits, and Climate Change

Louis Kaplow[*]

Abstract

This essay revisits the question of instrument choice for the regulation of externalities in the context of climate change. The central point is that the Pigouvian prescription to equate marginal control costs with the expected marginal benefits of damage reduction should guide the design of both carbon taxes and permit schemes. Because expected marginal damage rises nonlinearly, a corresponding nonlinear tax – or an equivalent price implemented through a quantity-adjusted permit scheme – is second best. Also considered are political factors, distinctive features of regulating a stock pollutant, and ex ante distortions due to the anticipation of transition relief (such as by receiving more free permits for greater emissions). Finally, distributive concerns are examined, with emphasis on the conceptual and practical benefits of addressing distributive issues with the tax and transfer system rather through adjustments to regulatory schemes that usually render them less effective.

JEL Classes D61, D62, H21, H23, K32, Q52, Q54, Q58

Taxes, Permits, and Climate Change
Louis Kaplow © 2010

1. Introduction

Kaplow and Shavell (2002) offer a systematic critique of the view associated with Weitzman (1974) that the choice between corrective taxes and quantity targets is highly contingent in a world of uncertainty.[1] They emphasize that this conventional understanding, heavily represented in texts and surveys on environmental economics as well as in more general textbooks, rests on a poorly appreciated and dubious set of assumptions.

First, in the Weitzman story, taxes are constrained to be linear even though marginal harm is taken to be nonlinear (typically, rising in the quantity of emissions), so the Weitzman tax instrument violates the basic Pigouvian prescription that the marginal tax rate equal marginal harm. Second, taxes are taken to be fixed for all time, even though any error in setting the (constrained-to-be-linear) tax rate would become immediately apparent to the government. Relaxing *either* assumption restores the superiority of taxes over quantity regulation. Moreover, neither assumption was well motivated by Weitzman or in the subsequent literature. Nonlinear taxes are simple to state in many settings, and tax rates (really prices) of all sorts – whether set by government or the market – are routinely adjusted in response to changed circumstances.

Second, Kaplow and Shavell (2002) extend their analysis to permit schemes. On one hand, this instrument has the familiar and important advantage over command-and-control quantity regulation that a given quantity target is achieved at minimum cost.[2] On the other hand, like familiar schemes, the quantity is set by fiat, once and for all. Kaplow and Shavell emphasize, however, that this latter defect is not inherent in permit schemes because, unlike command-and-control regulation, permit markets generate a price. Hence, the government learns whether firms' marginal costs (which in equilibrium equal the market price for permits) exceed or fall short of the marginal benefit of emissions reductions at the current targets. Drawing on prior literature, they describe a number of ways that the permit quantity could be adjusted to achieve the (same) second-best optimum produced by taxes (that are either explicitly nonlinear or adjustable). Relatedly, they point out that there is an underlying duality between properly designed tax and permit schemes: the former constitutes a quantity-dependent price, whereas the latter is a price-dependent quantity. In both cases, firms' information about their control costs is harnessed, so that each firm's marginal cost is equated to marginal benefits.

Weisbach (2010) devotes the first half of his fine paper to applying and elaborating Kaplow and Shavell's (2002) critique of Weitzman (1974) in the context of controlling greenhouse gas emissions. His motivation is that, sadly, just as the limitations of Weitzman were underappreciated a decade ago, the same holds true today, and such misunderstanding may lead

---

[1]In addition to Weitzman (1974), see also Adar and Griffin (1976), Fishelson (1976), Roberts and Spence (1976), and Rose-Ackerman (1973).

[2]This traditional advantage of permits (or taxes) is, unfortunately, getting far too little attention, notably, in the United States, where aggressive fuel economy standards, biofuels mandates, and other requirements may well be imposing huge costs relative to potential benefits, a waste in itself and a phenomenon that may reduce the political will to pursue further control through taxes or permit schemes.

to the design of regulatory schemes that are substantially and needlessly inefficient. Given the high stakes, such an outcome would be unfortunate for the United States and the rest of the world. In addition to economic waste, on a potentially vast scale, there is also the concern that one may control too little despite possibly catastrophic consequences.

This essay seeks mainly to reinforce and further extend the argument that the standard Weitzman assumptions should carry far less if any significant weight in instrument choice and design, particularly in the context of combating greenhouse gases. Section 2, by way of background, articulates Weitzman's (1974) assumptions and why they matter and then relates that framework to optimal policy. Section 3 analyzes nonlinear taxes, section 4 considers permit schemes, and section 5 brings the analysis of the two side by side to consider the underlying similarities, political appeal, and also some important complications in the climate control context. Section 6 notes two important additional considerations: the impact of uncertainty about future "prices" (whether tax rates or permit prices) on innovation and the importance of transition issues (notably, grandfathering based on prior emissions). Finally, section 7 addresses concerns about distribution, which are central in a number of the other papers in this volume as well.

2. Background

2.1. Weitzman Framework. – Weitzman (1974) compares quantity regulation and taxation in a world of uncertainty. Firms' marginal control costs are taken to be known by firms themselves but not observable by the government.[3] Realistically, we can think of this gap as capturing the government's residual uncertainty: that is, the government may have significant knowledge of control costs, but not as much as do individual firms. Regarding the control of greenhouse gas emissions, where the magnitude of marginal control costs in the future is of central importance, we are also concerned about the quality of firms' projections relative to those of the government.[4]

Marginal harm is likewise subject to uncertainty, a point of undoubted significance in the present setting. A further point, also quite pertinent, is that total marginal harm is not taken to be constant, but rather rising with the level of emissions, an uncontroversial and important factor with regard to climate change.

Weitzman's instruments are quite simple: a quantity regulation specifies the quantity of

---

[3]Strictly speaking, the Weitzman scenario envisions a single firm, and Kaplow and Shavell (2002) devote considerable attention to differences that may arise in considering settings with multiple firms. Given the nature of the climate change problem, the focus throughout this essay is on the case in which there are a large number of firms.

[4]One view would be that, at any point in time, the government's knowledge of the economy-wide marginal control cost would be as good as firms' knowledge. In this case, both linear taxes and permit schemes could hit the second-best price target perfectly, without any further adjustment. However, since marginal control costs (especially in the future) can depend on technological change, learning curve effects, and other subtle phenomena, significant asymmetric information may remain. Following the standard approach, the remainder of the discussion here assumes that the government is at an informational disadvantage, although it should be noted that much of what is stated would be applicable if the government's knowledge of private control costs was as good as firms'. (Note that with both taxes and permit schemes, all that is relevant is the economy-wide marginal control cost in equilibrium, not any particular firm's marginal control cost schedule.)

emissions that firms may produce, whereas a tax instrument (which he calls a price) specifies a single, fixed price per unit of emissions.

Given that the government is uncertain about both marginal control costs and marginal harm, it is easy to see that neither instrument will hit the optimum except by chance. The optimal quantity is that for which marginal control costs equal the marginal benefits of limiting emissions. If the government picks quantity, it may choose one that is too high or too low. Similarly, if it picks a tax rate, the quantity thereby induced (by firms that each equate their marginal costs to the tax rate) may be too high or too low. A priori, either approach might do better, and Weitzman shows that, under some particular assumptions, taxes are better when marginal control cost curves are steeper than the marginal harm schedule, and quantity regulation is preferable in the converse case.[5]

2.2. Optimal Policy. – It is generally understood that the Pigouvian ideal can be achieved by setting an emissions tax equal to marginal harm. Crucially, even if the government does not know firms' marginal control costs, this instrument is sufficient. Faced with a tax (price) equal to marginal external harm, the externality is internalized and firms' quantity decisions will be optimal. The fact that a tax instrument harnesses firms' information about control costs is their central virtue, as economists have long appreciated.

One complication is that the government may have uncertainty about marginal harm, a point that is of obvious relevance to climate change. Although many economists seem to believe (or act as if they believe) that uncertainty about harm tends to favor a quantity instrument, this view is mistaken. Another section of Kaplow and Shavell (2002) offers a simple proof that the second-best optimal scheme involves setting the tax rate equal to the *expected* marginal harm, and this result holds for all manner of distributions of uncertainty about marginal harm.[6]

It is worth noting, moreover, that this logic extends to any shape of the marginal harm schedule, including the possibility of thresholds or tipping points. If there was a tipping point at a known quantity level, the tax rate would jump infinitely at that point. More relevant for the present setting, consider the case of a stock pollutant and significant uncertainty about the location of any such tipping point – to which, one might add, no one really believes that there is a literal tipping point, at which another molecule will trigger catastrophe. In this scenario, marginal harm may be thought to rise more steeply as the size of the stock passes some given point. However, that point is uncertain and, moreover, there is uncertainty about the likelihood

---

[5]To suggest the intuition, suppose first that marginal harm was constant (so the marginal harm curve is horizontal); then a single, linear, fixed-for-all-time tax would achieve the second-best optimum. If instead the marginal harm curve is vertical (marginal harm goes from zero to infinity at a threshold quantity), then it is obvious that a quantity regime (setting the quantity at the threshold) would achieve the second-best optimum.

[6]As shown in Weitzman (1974) and Stavins (1996) in the setting with Weitzman's assumptions, correlation between marginal control costs and marginal harm affects optimal instrument choice. With a nonlinear tax, it remains true that it dominates quantity regulation, but the optimal nonlinear tax would no longer equal marginal harm. (Firms' quantity choices given a tax rate convey information on control costs, which, since they are now taken to be correlated with marginal harm, bear on the optimal marginal tax rate.) However, in many settings, this complication seems essentially irrelevant. Uncertainty about marginal harm derives from uncertainty in climate models, the effects of warming on agriculture, and myriad other factors that for the most part have little to do with control costs (e.g., the future cost of alternative energy sources or of carbon capture). For an important qualification, see the discussion of the second further complication in section 5.3.

that an additional unit today will lead the stock to cross such a threshold at a distant future time. Taken together, the expected marginal harm schedule will be a probability-weighted sum (integral) over myriad possible scenarios. Even with fairly sharp thresholds in some or many scenarios, the expected value of all such schedules will undoubtedly be a smooth expected marginal harm schedule.

Therefore, with uncertainty about harm and firms' control costs, the optimal scheme is a tax that imposes a charge on firms equal to the expected marginal harm from emissions. Because marginal harm is taken to be rising with the level of emissions, the optimal tax schedule is nonlinear, likewise increasing with quantity. As Weisbach properly emphasizes, with a stock pollutant as in the present context, this schedule is nevertheless locally pretty flat: a few percentage points more or less of emissions in a given month or year will (relatively) have a far smaller effect on the size of the present and future level of the stock and thus little effect on marginal harm and, accordingly, on the optimal marginal tax rate.

By contrast, charging a significantly wrong tax rate for extended periods of time will have high social costs. If the tax rate is much too low for a decade, there will be substantial buildup of greenhouse gases that will cause long-lasting harm. Even if one could control more later (and to a degree one can and should), the cost of doing so will be greater than it would have been if the proper tax rate had been charged earlier, inducing firms to undertake reductions for which the marginal control cost was below expected harm. Likewise, if the price is much to high for an extended period, costs significantly in excess of any marginal benefit will be incurred.

3. Nonlinear and Adjustable Taxes

The basic Pigouvian principle is that, ideally, the tax rate at any moment in time should equal the current best estimate of the expected marginal harm from emissions. With a slowly evolving stock pollutant like greenhouse gases, as just explained, the expected marginal harm from incremental emissions is not a rapidly moving target. Even if, say, emissions were to jump by many percentage points above expectation in a given year, the long-run impact on the level of the stock would be a very small fraction of the short-run change in the flow. Accordingly, even in the original formulation in Weitzman (1974), a single, constant, fixed tax rate will be nearly optimal over short time horizons (where *short* should be interpreted as a few years, not a few days).

However, the evolution of the stock over longer time periods exhibits considerably more uncertainty, due primarily to uncertainty in control costs. Current control costs may not be well understood; imposing a new regime designed to reduce emissions below the baseline path will move actors (both firms and consumers) to new portions of their control schedules, about which even less may be known; and technological change further contributes to uncertainty over longer periods. Accordingly, the degree of control in the future, in response to a given tax rate, could be substantially more or less than may have been anticipated. In turn, the marginal benefits of control could be less or more (respectively) than predicted, so the future optimal tax rate would be lower or higher, to a degree determined by the then-existing stock level and, of course, the marginal damage schedule. This is the reason that an optimal tax regime consists of a nonlinear schedule rather than a single tax rate, fixed for all time.

Weitzman (1974), as mentioned, restricts the tax regime to be linear, that is, for there to be a single, quantity-independent tax rate, and he does not allow adjustment of the tax rate in light of the observed quantities. (Note that the quantities must be observed to apply even a constant tax rate.) That is, the tax schedule is stipulated to be nonoptimal and to be immune from correction in light of the information that inevitably flows in. Kaplow and Shavell (2002) as well as Weisbach (2010) emphasize that Weitzman offers minimal practical justification for these assumptions and that subsequent literature does little more and, worse, fails to emphasize or sometimes even to mention the role that these restrictive assumptions play. This state of affairs is all the more surprising since most expositions purport to be generic, institution-independent analyses of the regulation of externalities. (Finally, it should be mentioned that it is hardly clear that Weitzman himself has a strong, broad belief in the applicability of the assumptions in his 1974 paper. Indeed, in a subsequent publication, Weitzman (1978), he unapologetically assumes that the pertinent tax is nonlinear, and proceeds to examine its implications in a different model.)

In the present context, a nonlinear schedule – itself an instrument less complicated than is sometimes supposed – would take an even simpler form. As explained, the need for a nonlinear schedule is to allow the marginal tax rate that firms' face in a given time period to reflect the level of marginal harm implied by the current level of the stock of greenhouse gases. Accordingly, it would not even be necessary for individual firms to face a nonlinear schedule. All that is required is for the current period's (fixed, linear) tax rate to be read off the previously announced marginal harm schedule.[7]

By analogy, consider the fact that some tax rates (or nonlinear tax schedules, such as with an income tax) may have predetermined paths: for example, a scheduled rate increase may be phased in over a number of years, according to a stated formula or table promulgated at the outset. Likewise, it is not uncommon for some taxes (again, notably, income taxes) to have automatic inflation adjustments of tax brackets. A given inflation index level from one year may mechanically determine the bracket (or exemption or other) levels for the next. What these examples have in common is that features of a tax system can change over time in essentially automatic ways based on predetermined intentions. Furthermore, these adjustments can be made a function of conditions as they unfold over time.

It is obvious that the described scheme – one with carbon tax rates, say, in a given year to be set based on stock levels observed as of the preceding year – is exceedingly simple administratively. Nor is it the case that nonlinear taxes would strike citizens and politicians as odd. Many have lived with nonlinear income taxes for their entire lives. There are other, familiar uses of nonlinear pricing as well. Quantity discounts (prices that fall with quantity) are commonplace in numerous settings, some utilities charge nonlinear rates, and the legal system

---

[7]Actually, nonlinear taxes with more complexity are themselves quite straightforward. The personal income tax in the United States and in many other jurisdictions applies a nonlinear schedule to each individual taxpayer, even individuals of highly limited sophistication. They might look up their tax obligations on a simple table (which does the computations for them) or go to the expense of purchasing tax preparation software or paying a tax preparer. Even if it were thought to be helpful, therefore, to employ nonlinear schedules for greenhouse gas emissions on a firm by firm basis, within a single period, it is hard to see why this challenge is thought to be other than trivial, much less prohibitive.

sometimes employs nonlinear sanctions (sentencing guidelines determine fines and prison terms in sometimes complex ways that may depend nonlinearity on, say, the magnitude of a theft or fraud, and speeding tickets sometimes involve nonlinear charges depending on the extent by which the driver exceeded the stated limit).

Furthermore, nonlinear taxes make intuitive sense in the context of regulating greenhouse gases. Consider the nonlinear schedule for tax rates that is stated in advance, but with a fixed (linear) tax imposed on all emissions in a given time period. Suppose it turns out that, in the previous period, the stock rose more quickly than the implicit benchmark rate. All would readily appreciate that there was some need to catch up. Hence, the imposition, according to the predetermined nonlinear tax schedule, of a correspondingly higher tax rate in the subsequent period would seem entirely natural. Indeed, the failure to undertake any self-correction would seem strange. Note that this simpleminded intuition is roughly in accord with the core logic of optimal Pigouvian taxation. Being behind in control of the stock means that the expected marginal damage caused by incremental subsequent emissions is higher than otherwise, making a higher tax optimal. The higher tax, in turn, induces firms to be willing to spend more at the margin to curtail emissions – up to the point that their marginal control costs are equated to the now-higher tax rate.

Likewise, if it turned out that the stock was below the expected level – the prior taxes resulted in more reductions than anticipated – we would be ahead of the game and it would seem natural to ease off somewhat. And, of course, in this scenario, emissions are less costly at the margin, so a lower tax rate, leading to somewhat less effort at control, is optimal. In sum, in both situations, the nonlinear prescription is not only quite simple administratively, but it also is not at all difficult to rationalize, even to the uninitiated.

In the foregoing analysis, no interim political decisions or international renegotiations are required at any point. All that has been described follows from an initial determination, which needs to set a marginal tax rate *schedule* – which is identical to the expected marginal harm schedule – as a function of the level of the stock at various future time periods. This schedule is automatically responsive to control costs being higher or lower than expected. As Kaplow and Shavell (2002) explain, there is built-in feedback with tax schemes. Levels must be observed in order to impose the tax in the first place (and, note, this information requirement is no different from what is required under a quantity scheme, like permits). And those levels provide exactly the feedback that is necessary to implement the nonlinear schedule. Accordingly, so-called "regulatory stickiness" is beside the point.

One can take the analogy to inflation adjustments in an income tax or welfare programs. If the adjustment is made on a discretionary basis, requiring new legislation, we might imagine all manner of political complications that could lead to something other than smooth adjustments. On the other hand, if we (as is now done in many settings) state an adjustment rule, the changes are made annually and almost effortlessly. There is no agency engaging in elaborate deliberations and no lobbying. A statistical agency publishes the index, and it is automatically used by low-level bureaucrats to compute the new schedules, which are then published. The notion that employing nonlinear schemes, or regularly adjusting linear tax rates, is subject to all sorts of political problems so as to render it highly impracticable or unreliable seems to be

without foundation.

There is, however, another important channel that would require some sort of deliberative process: new information about expected marginal harm. As explained in section 2.1, the optimal nonlinear tax schedule is one that is set equal to the expected marginal harm schedule. This approach achieves the second best. However, in the case of climate change, we anticipate that our knowledge of expected marginal harm will evolve, possibly quite significantly, over time as a result of new information and better models. It follows, therefore, that it would be optimal to change the schedule itself. For example, if we learn that warming is occurring more rapidly, it is optimal to impose higher taxes, whereas if ways to mitigate harm develop more quickly than anticipated, the tax schedule optimally would be lowered.

How this challenge should best be met is an important problem, and one worth more thought up front. Short of a full global renegotiation, it might be possible to delegate the task to an agency. In any case, when the control instrument is a tax schedule, the nature of the required adjustments are transparent and straightforward (that is, given the highly complex new information on marginal harm). Furthermore, it might be possible to build even some of this flexibility into an original schedule. For example, there could be a formula for how the schedule should be adjusted based on new findings about, say, how much warming is associated with various levels of greenhouse gas concentrations, so all the agency or negotiating group would have to determine would be that new relationship. Of course, such reassessments would naturally become the occasion for attempts to renegotiate a previous deal, a problem that seems inevitable under any approach to the problem of climate change.

4. Permits and Quantity Adjustment

As explained in section 2, it has long been understood that setting quantity targets tends to be inefficient when the regulator is uncertain about firms' control costs. Except by chance, the quantity will be set too high – when marginal control costs are less than expected marginal harm, the externality will be excessive relative to its optimal level – or it will be set too low – in which case non-cost-justified control will be undertaken. Moreover, as mentioned above, there may exist substantial uncertainty about firms' control costs, especially when setting quantity targets well into the future. If technological progress in reducing costs is more rapid than anticipated, valuable control opportunities will be forgone under a quantity target, whereas if projected breakthroughs do not occur, huge costs may be inefficiently borne.

This deficiency of quantity targets was traditionally associated with command-and-control regulation but is equally applicable to permit schemes set to meet hard targets. This point is equally true if those targets are changing over time, such as by becoming stricter over time in anticipation of cost reductions. It remains true that the target is determined by regulators' preconceived estimates of what firms' costs will be over time rather than being based on what those costs actually turn out to be.

Ever since Roberts and Spence's (1976) important paper, however, it has been appreciated by some theorists of environmental regulation that this need not be so. As Kaplow and Shavell (2002) explore more broadly, permit schemes have a highly attractive feature – in

addition to their well-known ability to minimize control costs across firms for a given quantity target. The key point is that the market price of permits not only coordinates control efforts efficiently, wherein each firm equates its marginal cost to the common market price, thus ensuring that all firms' marginal control costs are equated to each others', the condition for cost-minimization. As well, this price signals to the public regulatory authority the value of the current level of firms' marginal control costs. And this is precisely the piece of information that the regulator initially lacked, making it have to guess at the optimal quantity. (Note that command-and-control regulation, by contrast, does not yield such a price, and thus no matter how badly off is the quantity target, the regulator will not learn the magnitude or even the direction of its error from firms' behavior.)

Armed with knowledge of the permit price, the regulator can adjust the quantity of permits so as to achieve the second best, under which firms' marginal control costs are equated to expected marginal harm – just as is true under the nonlinear tax. Doing this is, in principle, straightforward: The regulator merely looks up the current quantity on the marginal harm schedule and sees what is the marginal harm at that quantity. With the nonlinear tax scheme, the regulator then announced that figure as the tax rate. With a permit scheme, the regulator merely has to compare this figure to the current permit price. If the price is below the tax rate, it optimally would be raised – and this is accomplished by reducing the quantity of permits. If the price is above the desired tax rate, it optimally would be lowered – by increasing the number of permits. Although the regulator may not hit the target precisely and immediately – i.e., the point at which the permit market price equals the ideal tax rate at the given quantity (here, the level of the stock) – it seems plausible that the regulator could come close fairly quickly.[8]

There are many methods of permit quantity adjustment, a number of which are described in prior theoretical papers (many of which are discussed in Kaplow and Shavell (2002) as well as in Weisbach (2010)). If permits were perpetual, one could sell additional ones or buy some back if the quantity was too low or too high (i.e., if the market equilibrium permit price was too high or too low).[9] More likely, permits would be issued and reissued annually or over some other period of moderate duration. In that case, the next period's issuance would be determined, as it would in any case, by some preexisting schedule or mechanism, and the foregoing analysis indicates that such a scheme optimally would adjust the number of permits in light of the previous equilibrium permit price so that the quantity going forward will more closely equate the

---

[8]Note further that permit schemes that adjust quantities to maintain target permit market prices are like nonlinear taxes in most other respects. For example, the central result of Strand (2010), in this volume, that exporters behave strategically differently, to importers' disadvantage, under permits compared to under taxes presumes that the permit regime has hard quantity targets. If instead price targets were employed, this difference (and disadvantage of permit schemes) would vanish.

[9]Alternatively, one could redesignate the units of existing permits. For example, instead of selling additional permits to raise the total by two percent, one might make prior permits worth (in terms of emissions or units of carbon) 1.02 times as much as before. A benefit (compared to sales and repurchases) is that, by leaving undiluted or uninflated the value of prior permits, permit prices would not be as directly influenced by the anticipation of subsequent expected permit purchases or issues. All perpetual schemes (and some others) face the challenges that firms with market power may have incentives to seek to manipulate permit prices to influence subsequent government policy. However, if permits are widely dispersed, as inevitably they may need to be, this problem may not arise. (A further complication is that a firm may seek to accumulate a large market share of permits, leasing their use to firms that actually need them; this practice might accordingly be prohibited or limited.)

market permit price with the expected marginal harm at the appropriate quantity.[10]

      Reflecting on such permit schemes with quantities that are adjusted in one manner or another over time, it seems clear that one can accomplish essentially the same outcome as with a second-best nonlinear tax schedule that equates the tax rate to expected marginal harm at the stated quantity. The up-front requirements are the same: This harm schedule must be determined, as best as can be done. Then, each period, instead of using the schedule to determine that period's tax rate, as a function of the observed quantity (projected forward to the next period), one would attempt to induce a market equilibrium permit price at the same level by selecting an appropriate quantity of permits.

      Such adjustments could be largely automatic and formulaic, rather than the product of discretionary, regular reconsideration. Also as with nonlinear taxes, if new information comes in regarding the expected marginal harm schedule, some reanalysis (such as by an agency delegated to perform such a process) or renegotiation would be required.

      Furthermore, note that these permit quantity adjustments should appear entirely natural to unsophisticated onlookers. If permit prices are much lower than expected, it will be clear that control costs are lower than anticipated, so it is sensible to lower the target quantity (that is, to aim for stronger emissions control), which is precisely what the permit quantity adjustment does in this situation. Likewise, unexpectedly high prices signal that costs are greater than imagined, calling for some relaxation of the target.

5. On the Comparison of Taxes and Permits

      5.1. Tax and Permit Duality. – One of the important observations in Kaplow and Shavell (2002) is that nonlinear taxes and adjustable-quantity permit schemes can be viewed as duals to each other. A nonlinear tax of the sort described in section 3 is one in which the price (tax rate) depends on the quantity. An adjustable permit scheme as described in section 4 is one in which the quantity depends on the price (in the permit market). Moreover, it is not simply that one is a quantity-dependant price and the other is a price-dependant quantity. It is further true that the second function is simply the inverse of the first function. It's the same curve, just flipping the *x* and *y* axes.

      In addition, this theoretical equivalence is borne out by the particular analysis of practical mechanics in sections 3 and 4. Nearly every feature examined in section 3, for nonlinear taxes, has a close analogue in section 4, for permits, and conversely. That is, the close conceptual linkage is also reflected to a great degree in administrative aspects of implementation. There are obviously limits to this latter claim, some of which are explored in Weisbach (2010), but the central point, in harmony with Weisbach's theme, is that there is far more in common than meets

---

[10]Another clever variation – first suggested by Roberts and Spence (1976, appendix) and also examined by Collinge and Oates (1982) and Laffont and Tirole (1996) – is to issue permits with different exercise prices, with those prices set to simulate the marginal damages schedule. The exercise price of the i^{th} permit would be the marginal harm of the i^{th} unit of emissions. In that way, the lowest numbered permits would be used first, permits would be used until the marginal permit in use had an exercise price just equal to firms' marginal costs in equilibrium, and each permit would trade for the exercise price of that marginal permit minus the particular permit's exercise price.

the eye. Indeed, for purposes of analysis, it is a good starting hypothesis that properly constructed tax and permit schemes are equivalent; differences need to be demonstrated, with an explanation of why features deemed unique to one of the approaches are not or cannot be replicated in the other. On reflection, given the underlying, abstract similarities – one might say near identity – this conclusion should not be surprising.

  5.2. Political Appeal. – This essay and this author's expertise do not extend to political considerations. However, some brief observations are in order. First, as already explained with regard to both nonlinear taxes and adjustable-quantity permit schemes, the nature of the price or quantity adjustments over time is in strong accord with common sense. Hence, there is not an obvious problem with technical economics producing a complex, subtle, or counter-intuitive policy prescription that must be somehow explained or snuck past legislators and citizens.

  Second, there seems to exist some direct appeal to certain factions in the climate control debate, and these sorts of control schemes may actually facilitate compromise in light of actual (or feigned) disagreements on important facts. For example, strong environmentalists often advance the view that control costs – especially in the long run, once a serious regime becomes a reality, making technological advance highly profitable – will become much lower than fear-mongers assert. If these proponents really believe this to be true, then nonlinear taxes or adjustable-quantity permit schemes should be highly attractive. If costs indeed will be low, the resulting level of control will be high – in particular, higher than would result from pre-established quantity targets that reflected a compromise between differing views (rather than the environmentalists' views being vindicated).

  Likewise, to those worried about the effect of stringent controls on the standard of living – either for rich or for developing countries – both nonlinear taxes and adjustable-quantity permit schemes have a built-in safety valve. If control costs do remain high, the stringency of the controls automatically loosens.

  Notice that there is a close connection between these observations and the original theoretical justification for using (nonlinear) taxes rather than hard quantity targets (originally understood to be implemented by command-and-control regulation). The problem with the latter is that the government learns nothing about firms' control costs and thus is unable to adjust targets accordingly over time. The beauty of taxes is that they harness firms' information about control costs. Permit schemes, through the market equilibrium permit price, do the same. Both types of systems – nonlinear taxes and adjustable-quantity permit schemes – use this information, implicitly or explicitly, to adjust quantity targets in light of the information thereby revealed, producing the second-best optimum (or a reasonable approximation thereof).

  Uncertainty about control costs, which becomes greater the longer one projects into the future, is not only a problem for a benevolent, all-powerful regulator, but also serves as an obstacle to reaching political consensus. It appears that the magic of nonlinear pricing, through taxes or adjustable-quantity permit schemes, may allow us to elide some of the political problem as well. Opponents in the debate about the likely levels of future control costs can agree to disagree. All that needs to be set is the expected marginal harm schedule, and the market does the rest. Of course, there is also substantial disagreement (to say the least) about future harm.

Accordingly, these schemes are no political panacea. However, given the huge magnitude of the political challenges in achieving consensus, the process can use all the help it can get.

5.3. Complications.[11] – One complication, the largest political challenge, concerns the inevitably international scope of the climate change problem. The coordination (free-rider) problem is massive. As a consequence, the task is even more complex and of more doubtful efficacy if constant renegotiation is required. This difficulty is one reason why it is thought to be important to strike a deal that reaches far into the future. However, a major challenge in so doing is that the future is highly uncertain, so any long-term deal is likely to become substantially suboptimal as time unfolds, above and beyond sacrifices to political expediency.

Nonlinear taxes and adjustable-quantity permit schemes address this problem in part. A central need for adjustment concerns ex ante uncertainty about future control costs, where more information will be revealed by marketplace activity over time. As already explained, these schemes have built-in, automatically adjusting flexibility in this regard. Thus, their efficiency advantage may also be a political advantage in yet another respect. However, as noted, this benefit does not extend to new information about marginal expected harm. To some extent, one might build some such information, perhaps as interpreted by a pre-designated expert agency, into the nonlinear tax or permit price schedule, but there are probably political limits on the extent to which this can be done.

A second complication concerns the stock nature of the greenhouse gas problem. As already noted, this is in one important respect an advantage: Because the stock changes very slowly relative to possibly unexpected short-term fluctuations in the flow, the optimal tax rate or permit price also changes very slowly, so there is little efficiency loss in a regime that adjusts taxes or permit quantities even with a substantial lag. (Fairly rapid adjustment would be feasible; the point is that, in the present setting, it does not seem to be very important.)

One problem, however, is that the stock nature of the pollutant greatly complicates the problem of estimating the marginal harm from a given current unit of emissions. In addition to all of the problems of climate change modeling and other aspects of ascertaining marginal damages, there is the point that the marginal harm of a current unit of emissions depends in significant part on estimates of future marginal control costs (an interdependency that is absent with a flow pollutant). For example, if future costs turn out to be lower than anticipated, the self-adjusting system will generate a lower future stock, and this in turn will reduce the expected marginal harm of *current* emissions. Therefore, estimates of contemporaneous expected marginal harm will in part reflect expectations about future control costs. Despite this complication, it remains true that the optimal present tax rate or permit price equals the best estimate of expected marginal harm (taking this interdependency into account). The current point is that determining this magnitude will be even harder than the daunting task one envisions considering only the harm side of the calculus.

6. Additional Considerations

6.1. Future Price Uncertainty and Innovation. – It is sometimes feared that uncertainty

---

[11]Obviously, there are myriad additional complications, some of which are addressed in Weisbach (2010). This essay examines two that are closely related to the foregoing discussion.

about future prices, whether future tax rates under a nonlinear tax scheme or future permit prices under an adjustable-quantity permit system, will inhibit innovation (and otherwise impose social costs). As a first approximation, it seems that this is a non-issue in the following sense. Markets generate investments based on expected future payoffs. If there is genuine future risk – climate change has a substantial enough magnitude to contribute to systematic risk, which even in theory cannot be fully diversified – ex ante investments will take this uncertainty into account. However, that investments reflect future uncertainty tends to be socially optimal. Market risks are real. They cannot be wished away by government or anyone else. And, given their presence, decisions should be optimized in light thereof.[12]

The situation with regard to investments pertaining to climate change seems like that for all manner of investments in a market economy.[13] Especially in recent decades, where massive investments have been undertaken in technology sectors – computer chips, communications, the Internet, software – it is clear that the market operates tolerably well in light of massive uncertainty. The market is hardly perfect. One particular concern is innovation spillovers, which cannot be fully captured and thus lead to inadequate investment, particularly in basic research. Even optimal and well-enforced intellectual property rights regimes involve distortion, and determining the optimal regime is no easy task. Accordingly, there may be a substantial role for government support of basic research and for international cooperation along this dimension as well. However, just as it probably would have been a poor (if not disastrous) alternative to have the government pick winning and losing technologies in high-tech industries – and substantially problematic for the government, say, to peg computer chip or broadband prices in the distant future – so too government optimally would not interfere with permit market prices, that is, except to equate them to best estimates of expected marginal harm. Likewise, tax rates optimally would be set with regard to the estimated magnitude of the externality, not tinkered with in order to provide a more stable investment environment.

The importance of this point can be made more concrete by considering the actual impact of future price uncertainty. Suppose, for example, that there is a ten percent probability that future prices will be much higher than standard estimates, perhaps because it is possible that certain cost-reducing technologies will confront unanticipated hurdles. In that case, it is important for investors to be unsheltered from this risk. The prospect of such high prices will induce greater ex ante efforts toward alternative means of reducing emissions, whether through the development of alternative technologies or more aggressive conservation. Similarly, if there is a serious prospect that prices might be much lower than consensus predictions, perhaps due to the possibility of major breakthroughs, then the dampening effect of this prospect, which is to reduce the incentive to pursue other means of emissions control, is optimal to that extent. Granted, there can be bubbles and other phenomena that render markets (sometimes highly) imperfect guides, and it is desirable to minimize such pricing defects where possible. However, a general policy of arbitrarily dampening the range of tax or permit price movement is a blunt

---

[12]There will inevitably be mechanisms to allocate risks to some degree – investors hold diversified portfolios; futures markets may allow some risks to be hedged – but whatever is the residual risk, it reflects real uncertainty about economic activity and hence should be taken into account by private actors.

[13]The analysis is closely related to the core argument against transition relief (considered in the next subsection) that its anticipation distorts ex ante investment incentives. See, for example, Kaplow (1986, 1992) and Shaviro (2000).

instrument that would likely cause avoidable inefficiencies. Even worse would be to move more toward command-and-control alternatives that eschew the benefits of the price mechanism to an even greater degree in the misguided pursuit of predictability.

   6.2. Transition, Grandfathering, and Ex Ante Incentives. – The transition problem is substantial for the regulation of greenhouse gases, and this issue, viewed generally, has received sustained attention in some of my own prior work (e.g., Kaplow 1986, 1992, 2003, 2008b) and that of other scholars (e.g., Shaviro 2000). For the present, it is useful to emphasize one teaching: the importance of accounting for ex ante investment incentives, particularly when a substantial transition is anticipated for a long period of time. With climate change, it is likely that major agreements still may be many years off. And the prospect of a substantial, global regulatory regime has been under serious consideration for an extended period of time. Moreover, whatever national legislation may be enacted in the immediate future or global agreement may be ratified in the near term, it seems likely that significant alterations will be ever on the horizon, making ex ante concerns relevant well into the future.

   If firms (and countries) anticipate substantial grandfathering, great inefficiency can arise. One particular concern is misdirected investment that is counterproductive in combating climate change. One form of grandfathering would involve distributing free permits based on pre-enactment emissions. Another would be granting analogous preferences under a tax scheme (for example, exemptions for pre-enactment quantities of emissions might be provided).[14] To the extent that any such relief is based on prior investment, firms and countries have inadequate ex ante incentives to begin bringing their investments and policies into line with the anticipated new regime, and they may even behave counterproductively in order to secure greater preferences upon enactment. As just one illustration, if such grandfathering is long anticipated, China may build thousands more dirty coal plants than they would if their permit allotment had to be purchased or was distributed, say, based on their population or some other factors – that is, factors independent of how much annual greenhouse gases they manage to be pumping into the atmosphere before a deal is struck.

   Granting free permits or otherwise providing relief to costly, pre-existing sources is often viewed as a tradeoff between equity and political feasibility. However, the world may well be currently incurring a tremendous cost by following a course in which firms and countries expect to be rewarded for increasing rather than reducing emissions in the interim – and this process has been going on for over a decade and may continue for quite some time. Likewise, if subsequent renegotiations are expected to favor those who have done the worst job complying with a new regime, any agreement is much less likely to be successful going forward. Accordingly, any moves that can be made now to signal credibly that such transition relief will not be provided (or will be less generous than currently anticipated) could have a substantial social payoff – the more, the sooner, the better.

---

   [14]A difference would arise if the free permits were tradeable whereas the tax preferences were not (i.e., firms could not directly or indirectly transfer their units of exempt emissions to other firms with greater control costs). On one hand, free transferability ex post is more efficient since those receiving the relief may have lower marginal control costs than do other firms. On the other hand, the prospect of transferability, by raising the value of relief, worsens the ex ante incentives discussed in the text.

7. Distribution

Both tax and permit schemes raise serious concerns – normative and political – about their impact on the overall distribution of income.[15] The subject comes up in Weisbach (2010), receives greater attention in a number of other papers in this volume, notably Burtraw et al. (2010) and Rausch et al. (2010), and plays an important implicit role in de Gorter and Just (2010) and Parry (2010). A number of powerful lessons are important to keep in mind.

First, it is often supposed that one must design the climate change interventions themselves to meet distributive objectives. For example, free permits or other leeway might be given to the poor or to people who drive in order to mitigate the regressive distributive incidence of gasoline price increases. This view, however, is fundamentally mistaken and tends to confuse analysis and generate policies that are less efficient than available alternatives. The most straightforward and (approximately) best approach is to design an optimal environmental control scheme, fully allowing price increases, such as in the price of gasoline, to occur in order to induce all manner of efficient behavior (driving less, purchasing more fuel-efficient vehicles), and then achieve distributive objectives through the income tax and transfer system. That is, rather than contorting environmental policy, we can implement it efficiently and adjust the general distributive apparatus of the economy to provide any desired offsets.

One particular means of doing so, which is a central focus of my recent book, Kaplow (2008a), as well as of my writing on optimal environmental policy and regulation specifically (Kaplow 2004, 2006), is to adjust the income tax and transfer system so that the combined scheme is distribution neutral. For example, if a carbon tax or permit scheme raises the price of gasoline, consumed disproportionately by the poor, one might lower taxes (or raise an earned income tax credit and/or adjust other transfer programs) to make up the difference. If environmental benefits are disproportionately concentrated on the rich, their income taxes might be raised in an offsetting manner.

This simple prescription follows from familiar economic principles: It is generally best to use a separate instrument to address each distinct problem; moreover, for each problem, it tends to be desirable to employ that instrument which addresses it most directly. Use emissions taxes or permit schemes to get the carbon price right (which will induce efficient behavior on all margins related to the environmental challenge) and use the income tax and transfer scheme to produce the desired distribution. (The distribution-neutral approach need not be optimal, of course, since the pre-existing distribution need not be optimal, but its availability suggests the ability to separate the two problems.)

---

[15]There are other dimensions of distribution, such as geographic incidence, that may have political relevance but are set to the side here. See Burtraw et al. (2010). Note, however, that protecting the preexisting distribution on such dimensions based on pre-enactment emissions patterns can both produce the sorts of inefficiencies discussed in section 6.2 and also (if design provides long-term, ongoing compensation) perpetuate or even increase inefficiencies in subsequent behavior. For example, if energy usage is much higher in some climates, and in light of the greenhouse gas problem the true total social cost of energy is much higher than previously thought, then a part of long-run efficient adjustment involves a different geographic distribution of the population. Over the pertinent time horizons in climate change debates (covering decades, or even a century or more), margins of this sort can be quite significant: Population distribution will be much different in 2050 than it is today, and it is important that the incentives guiding such movement reflect the true social costs of living in different regions.

Another often overlooked feature of distribution-neutral implementation is that it not only neutralizes distributive effects (by definition) but also, as a crude first approximation, is neutral with regard to labor supply. (The assumption that generates this result is that labor is weakly separable in utility from other features of the problem, notably, different forms of consumption and environmental effects.[16]) The basic intuition is that, with distribution-neutral implementation, there is the same marginal return to earning more income after the reform as there was before, all things considered, so individuals would be induced to choose the same level of labor effort.

These ideas are also pertinent to the common practice of comparing the implications of different means of implementation regarding the collection and distribution of revenues. Auctioned permits raise revenues compared to free distribution; likewise, tax schemes add funds to the treasury. Differences in revenue raised are thus seen by some to be important distinctions among permit schemes and between permits and taxes. Moreover, it is common for economic analyses to compare different means of distributing these collections: Revenue might be rebated lump sum (equal per capita), by proportionally reducing labor income taxes, or otherwise. Indeed, some of the aforementioned papers in this volume devote substantial attention to these comparisons.

However, in light of the foregoing discussion, one can see that such analyses are often incomplete and, as a consequence, misleading. Sometimes different revenue treatments (on the collection and the disbursement side) are compared in terms of their efficiency costs, without regard to distribution. However, the differential efficiency cost of different revenue treatments is largely or fully a consequence of differences in redistribution, as emphasized in Kaplow (2004, 2006, 2008a). Should we be surprised that distributing funds lump sum has a higher efficiency cost than using them to reduce marginal tax rates? After all, comparing these two approaches, the environmental intervention is held constant, and the only difference is that the former, by comparison to the latter, has higher marginal tax rates to finance a more generous lump-sum transfer. This is a quintessential increase in redistribution through the tax system which really has nothing to do with regulating greenhouse gases. And should we be surprised that this redistributive difference is associated with greater overall distortion? As explained, a distribution-neutral implementation would wash out this efficiency consequence (the distortionary effect of redistribution on labor supply) and the distributive difference, leaving just the intrinsic effects of the environmental regulation under consideration.

Some other analyses do attend to distributive effects but, as mentioned, tend to see them as tied to particular means of deploying various regulatory instruments. Since, as explained, one can adjust the tax and transfer system so as to achieve distribution neutrality – or any other desired distribution, for that matter – distribution (and labor supply distortion associated therewith) should be seen as a largely separable problem from that of addressing climate change. Moreover, if academic analysis is to usefully and clearly guide the political process, which is

---

[16]For formal analysis, see Kaplow (2006). If, instead, gasoline is, say, a leisure complement, as indicated by West and Williams (2007), then it is optimal to set the corrective tax wedge on gasoline above the Pigouvian level (because doing so will further reduce gasoline use, which under the maintained assumption will also reduce leisure, which is distorted upward by the income tax and transfer system).

concerned with distributive matters, it seems best to present versions of implementation that are explicitly distribution neutral – to which analysts might append other distributive scenarios, clearly described, with accompanying distributional tables or diagrams. The distribution-neutral case is a useful analytical and political benchmark that facilitates communication and clarifies thinking.

Observe further that there is a potentially important political advantage of distribution-neutral policy packages. Part of the difficulty of reaching agreement, whether within a single country's legislature or in an international multilateral bargain, concerns such distributive effects.[17] Indeed, they may be one of the greatest obstacles. Note, however, that a feature of an efficient policy change that is implemented in a distribution-neutral fashion is that it tends to result in a Pareto improvement. Enhanced efficiency means a larger pie, and distribution-neutrality means that this expanded pie is divided among parties the same way that it was before. Therefore, everyone's slice is larger. Although this depiction of political deal-making is highly oversimplified, it seems plausible that this factor should carry some weight.

8. Conclusion

The principal conclusion of this essay is that there exists a strong duality between nonlinear tax schemes and adjustable-quantity permit schemes. Moreover, both possess a significant advantage over fixed-quantity permit schemes because the level of control of greenhouse gases under the latter could readily be much too great or too lenient, especially in the long run. The Pigouvian dictum that polluters should face a tax or price equal to the marginal expected damage that they cause is the appropriate benchmark for thinking about policies that regulate externalities, including greenhouse gas emissions.

Weitzman's (1974) framework, which compares a fixed quantity regulation to a fixed, linear tax scheme is substantially misleading and inapt with regard to climate control. In regulating a stock pollutant, the level of which evolves quite slowly, the second-best optimal nonlinear tax can be very closely approximated in administratively straightforward ways using both taxes and permits. Moreover, the adjustments over time that are necessary in light of the resolution of uncertainty about control costs – which is substantial when considering projections decades into the future – can be done in an essentially automatic, nondiscretionary manner. Unfortunately, any type of scheme will need discretionary adjustment, and possibly political renegotiation, as more is learned about the level of expected social harm due to greenhouse gases.

Furthermore, the built-in adjustment process in a nonlinear scheme, whether implemented through taxes or permits, is fairly intuitive and thus does not in itself present an obvious political barrier. Indeed, as explained, these adjustments seem so natural that it actually may appear more

---

[17]In the domestic context, the income tax and transfer programs are the core instruments to adjust distribution. Internationally, using similar logic, it tends to be best to use explicit transfer payments precisely because indirect means tend to involve additional, avoidable inefficiencies. In the present setting, it is familiar that one of the greatest sources of efficiency from an optimal international agreement is the potential to achieve cost-minimization across countries, with more control taking place in often poorer countries where opportunities for cheap reductions may be large, which effort may be induced by appropriate side-payments.

surprising not to make them (such as would be the case under fixed-rate tax or hard-target quantity schemes).

Some additional considerations were also addressed. One often overlooked problem concerns transition, in particular that giving away permits or providing tax exclusions on the basis of existing emission levels provides perverse incentives for control prior to enactment. This problem is likely to be substantial in light of the fact that truly tough limitations are unlikely to become binding for quite some time, making it far more attractive for both individual sources within countries and many nations to expand activities that will cause additional interim damage and make more daunting the achievement of future emissions reductions.

Another dimension – addressed by many of the papers in this volume, but often analyzed incompletely and misleadingly – concerns distributive effects, particularly with regard to the overall distribution of income within and between countries. The central point to keep in mind is that the tax and transfer system (including transfers between countries) typically offers the most efficient means of addressing distributive concerns. In particular, distribution-neutral implementation – combining an efficient greenhouse gas regulatory scheme with tax and transfer adjustments that, taken together, keep the overall distribution constant – is quite appealing as an analytical construct. It clarifies thinking, for example by distinguishing distortions due to increased redistribution from distortions related to externality regulation. Furthermore, distribution-neutral systems may be easier to communicate to policy-makers (who could also have non-distribution-neutral proposals compared side-by-side to clarify distributive effects), and they also have some pragmatic appeal, for efficient (pie-expanding) regulatory schemes that are also distribution-neutral will make all groups better off.

In closing, it is useful to emphasize that basic economic principles – like setting prices for externalities that equal marginal harm, and carefully distinguishing the task of externality correction from that of redistribution – are important to keep in mind and to disseminate, even amidst ongoing, intense political debate and negotiation about practical problems. In the short run, better policies can be designed and implemented, whether they take the form of taxes or permit schemes.

Also important is the long run, especially since the greenhouse gas problem, unfortunately, is likely to be with us for generations, and further regulation and modification are inevitable in the mid and distant future. Many will recall that, only a few decades ago, economists' advocacy of permit schemes as superior to command-and-control regulation was often viewed as evil because such an approach entailed approval of paying to kill people. Had proponents of permits given up the battle then, existing environmental regulation would almost certainly be more costly and less effective, and the prospects for addressing climate change would be even more grim.[18]

Pushing the point that optimal control is cost-effective control, which entails not only cost minimization in achieving a given target but the choice (and modification) of targets that reflect an equation of marginal control costs and marginal benefits, both of which may change

---

[18]Even after the introduction of the first SO_2 trading system in the early 1990s, media responses (including a prominent cartoon) portrayed permits as legalized certificates to kill people.

substantially over time from current estimates, is of first-order importance. No final, definitive, international regime for addressing greenhouse gases will be in place any time soon, and whatever arrangements are adopted, there will be continuing pressure to modify them. Hence, even ideas that cannot fully and persuasively be presented in the short run are worth advancing for the long run.

References

Adar, Zvi, and James M. Griffin. 1976. "Uncertainty and the Choice of Pollution Control Instruments," 3 *Journal of Environmental Economics and Management* 178–88.

Burtraw, Dallas, Margaret Walls, and Joshua Blonz. 2010. "Distributional Impacts of Carbon Pricing Policies in the Electricity Sector." In Gilbert E. Metcalf, ed., U.S. Energy Taxes.

Collinge, Robert A., and Wallace E. Oates. 1982. "Efficiency in Pollution Control in the Short and Long Runs: A System of Rental Emission Permits," 15 *Canadian Journal of Economics* 346–54.

de Gorter, Harry, and David R. Just. 2010. "The Social Costs and Benefits of U.S. Biofuel Policies with Pre-existing Distortions." In Gilbert E. Metcalf, ed., U.S. Energy Taxes.

Fishelson, Gideon. 1976. "Emission Control Policies under Uncertainty," 3 *Journal of Environmental Economics and Management* 189–98.

Kaplow, Louis. 1992. "Government Relief for Risk Associated with Government Action," 94 *Scandinavian Journal of Economics* 525–41.

_____. 2003. "Transition Policy: A Conceptual Framework," 13 *Journal of Contemporary Legal Issues* 161–209.

_____. 2004. "On the (Ir)relevance of Distribution and Labor Supply Distortion to Government Policy," 18(4) *Journal of Economic Perspectives* 159–75.

_____. 2006. "Optimal Control of Externalities in the Presence of Income Taxation," National Bureau of Economic Research Working Paper 12339.

_____. 2008a. *The Theory of Taxation and Public Economics*. Princeton: Princeton University Press.

_____. 2008b. "Capital Levies and Transition to a Consumption Tax." In Alan Auerbach and Daniel Shaviro, eds., *Institutional Foundations of Public Finance: Economic and Legal Perspectives*, pp. 112–46. Cambridge: Harvard University Press.

_____, and Steven Shavell. 2002. "On the Superiority of Corrective Taxes to Quantity Regulation," 4 *American Law and Economics Review* 1–17.

Laffont, Jean-Jacques, and Jean Tirole. 1996. "Pollution Permits and Environmental Innovation," 62 *Journal of Public Economics* 127–40.

Parry, Ian W.H. 2010. "Tine to Raise Gasoline and Direct Taxes?" In Gilbert E. Metcalf, ed., U.S. Energy Taxes.

Rausch, Sebastian, Gilbert E. Metcalf, John M. Reilly, and Sergey Paltsev. 2010. "Distributional Impacts of a U.S. Greenhouse Gas Policy: A General Equilibrium Analysis of Carbon Pricing." In Gilbert E. Metcalf, ed., U.S. Energy Taxes.

Roberts, Marc J., and Michael Spence. 1976. "Effluent Charges and Licenses Under Uncertainty," 5 *Journal of Public Economics* 193–208.

Rose-Ackerman, Susan. 1973. "Effluent Charges: A Critique," 6 *Canadian Journal of Economics* 512–27.

Shaviro, Daniel. 2000. *When Rules Change: An Economic and Political Analysis of Transition Relief and Retroactivity*. Chicago: University of Chicago Press.

Stavins, Robert N. 1996. "Correlated Uncertainty and Policy Instrument Choice," 30 *Journal of Environmental Economics and Management* 218–32.

Strand, Jon. 2010. "Tax Caps as Climate Policy Instruments with Domestic and Imported Fuels." In Gilbert E. Metcalf, ed., U.S. Energy Taxes.

Weisbach, David. 2010. "Instrument Choice Is Instrument Design." In Gilbert E. Metcalf, ed.,

U.S. Energy Taxes.

Weitzman, Martin L. 1974. "Prices *vs.* Quantities," 41 *Review of Economic Studies* 477–91.

_____. 1978. "Optimal Rewards for Economic Regulation," 68 *American Economic Review* 683–91.

West, Sarah, and Roberton C. Williams. 2007. "Optimal Taxation and Cross-Price Effects on labor Supply: Estimates of the Optimal Gas Tax," 91 *Journal of Public Economics* 593–617.