

Hormesis, hotspots and emissions trading

Jonathan B Wiener*

Law School, Nicholas School, Sanford Institute, and Center for Environmental Solutions, Duke University, Box 90360, Durham, NC 27708-0360, USA

Instrument choice – the comparison of technology standards, performance standards, taxes and tradable permits – has been a major topic in environmental law and environmental economics. Most analyses assume that emissions and health effects are positively and linearly related. If they are not, this complicates the instrument choice analysis. This article analyses the effects of a nonlinear dose–response function on instrument choice. In particular, it examines the effects of hormesis (high-dose harm but low-dose benefit) on the choice between fixed performance standards and tradable emissions permits. First, the article distinguishes the effects of hormesis from the effects of local emissions. Hormesis is an attribute of the dose–response or exposure–response relationship. Hotspots are an attribute of the emissions–exposure relationship. Some pollutants may be hormetic and cause local emissions–exposure effects; others may be hormetic without causing local emissions–exposure effects. It is only the local exposure effects of emissions that pose a problem for emissions trading. Secondly, the article shows that the conditions under which emissions trading would perform less well or even perversely under hormesis, depend on how stringent a

level of protection is set. Only when the regulatory standard is set at the nadir of the hormetic curve would emissions trading be seriously perverse (assuming other restrictive conditions as well), and such a standard is unlikely. Moreover, the benefits of the overall programme may justify the risk of small perverse effects around this nadir. Thirdly, the article argues that hotspots can be of concern for two distinct reasons, harmfulness and fairness. Lastly, the paper argues that the solution to these problems may not be to abandon market-based incentive instruments and their cost-effectiveness gains, but to improve them further by moving from emissions trading and emissions taxes to risk trading and risk taxes. In short, the article argues that hormesis does not pose a general obstacle to emissions trading or emissions taxes, but that in those cases where hormesis does pose such a problem, a shift toward risk trading or risk taxes would be the superior route. *Human & Experimental Toxicology* (2004) 23, 289–301

Key words: hormesis; hotspots; instrument choice; emissions trading; risk; dose–response

Instrument choice has been a major topic of environmental law and environmental economics. But to date we have not addressed whether instrument choice should be sensitive to the particular toxicology of each pollution problem. How does the dose–response relationship for a pollutant affect the optimal choice of policy instrument to regulate that pollutant? James Hammitt's insightful article^{1,2} analyses many of the policy implications of a hormetic or J-shaped dose–response function (in which a substance is harmful at high doses but beneficial at low doses).^{2,3} Examples of hormesis may include vitamins, essential minerals, sunshine, red wine, oxygen, salt and perhaps many other substances, possibly even radiation and some pollutants. Most of Hammitt's article deals with the implications of hormesis for standard setting: how much protection against exposure the government should require. In the last section of the article,

Hammitt briefly addresses the implications of hormesis for instrument choice. Given a standard to be achieved, should the government use command-and-control instruments such as technology mandates, or fixed limits on the quantity of emissions at each source, or market-based incentive instruments such as emissions taxes or tradable emissions allowances? And, how is that choice affected by the shape of the dose–response function?

Economic theory and experience show that if the cost of emissions abatement varies across sources, emissions taxes and tradable emissions allowances will generally achieve a given standard at less cost (or achieve greater protection at a given cost) than will technology mandates or fixed emissions limits.⁴ Hammitt argues, however, that tradable emissions allowances are likely to have less of an advantage over alternative policy instruments if the dose–response function is hormetic than if it is linear. He says: 'Under the hormetic model, a simple economic incentive mechanism having a single tax or single type of tradable permit would not be anticipated to work as well as under the linear model'.¹

*Correspondence: Jonathan B Wiener, Law School, Nicholas School, Sanford Institute, and Center for Environmental Solutions, Duke University, Box 90360, Durham, NC 27708-0360, USA
E-mail: wienjer@law.duke.edu

Policy instruments and dose–response

Market-based incentive instruments such as emissions taxes and tradable emissions allowances reduce the cost of achieving a given level of protection (or achieve greater protection at a given cost) by allowing the degree of abatement to vary across sources. This variation is often termed intersource flexibility, ‘where flexibility’ or locational flexibility. In contrast, fixed emissions limits allow each source flexibility in its internal abatement or compliance method (intrasource or ‘how’ flexibility) (as do tradable emissions allowances), but no flexibility to vary emissions across sources. (Fixed emissions limits can allow ‘where’ flexibility within the source; this is typically called a plantwide ‘bubble’ approach to aggregate emissions from the source. Tradable emissions allowances expand the bubble to encompass more than one plant or firm across a region, thereby expanding the scope for intersource ‘where’ flexibility.) Technology mandates dictate specific abatement methods at each source and thus allow neither ‘where’ flexibility nor ‘how’ flexibility. If the cost of emissions abatement varies across sources, then ‘where’ flexibility enables more abatement to be undertaken at sources where it is less costly, and less abatement to be undertaken at sources where it is more costly. The result is the same aggregate level of abatement, but at lower total cost, and with some heterogeneity in the ultimate level of emissions across sources.

Consider a hypothetical lake with the pollutant smox emitted by four sources evenly spaced around its shore: A, B, C and D. As shown in Table 1, assume that without any required controls, each source emits 20 units of smox, for a total of 80. Now the government imposes limits to achieve a 50% aggregate reduction to a new total of 40 units of smox. The cost and location of this abatement depends on the policy instrument the government chooses. If the government imposes a technology mandate (smox scrubbers) predicted to cut emissions by 50% (to 10 units at each source for a new total of 40), the cost per source will be \$100, for a

total cost of \$400. If, however, the government imposes fixed emissions limits of 10 units per source, each will again emit 10 for a total of 40; source A will install the smox scrubber at a cost of \$100, but sources B and C will change their materials inputs to reduce emissions at a cost of \$50 each, and source D will change its process method to reduce emissions at a cost of only \$20, so the total cost will be \$220. ‘How’ flexibility saves \$180. If, alternatively, the government imposes tradable emissions allowances, issuing 10 allowances to each source, then again the total emitted will be 40. But now the amount emitted at each source may vary from 10. Source A will seek to avoid the costly smox scrubber and instead purchase 10 allowances from the other sources at some cost less than the scrubber’s \$100. Source D will seek to sell additional abatement at a price above its low cost of abatement (say, \$25 for the next 10 units). If A buys 10 allowances from D, then A emits 20, D emits zero, and B and C emit 10 each, for a total of 40. A will pay somewhere between \$26 and \$99 to D to purchase D’s 10 allowances; assume A pays \$30 to D, saving A \$70 while earning D \$5. B and C will each spend \$50 as before. The total abatement cost will be \$145 (\$50 each at B and C, plus \$20 at D for the first 10 abated and \$25 for the second 10 abated; the extra \$5 paid by A to D is a transfer). ‘Where’ flexibility saves \$75 compared to fixed emission limits and \$255 compared to technology mandates. A similar result would be obtained if the government imposed an emissions tax; source A would pay the tax while source D abates to avoid the tax and sources B and C abate partially and pay partially.

The example above omits the administrative costs of emissions monitoring and enforcement, which would be necessary under all three policy instruments. It also omits risk–risk tradeoffs with emissions of other pollutants or into other media. It also shows A purchasing allowances to return to 20 units of emissions and not higher, on the intuition that if A’s unregulated emissions were 20, A would not now find it more profitable to emit more than 20 (at the cost of additional allowances or tax payments);

Table 1 Illustrative emissions and costs under alternative policy instruments

Source	No regulation		Technology		Fixed emissions		Tradable emissions	
	Smox emitted	Abatement cost	Smox emitted	Abatement cost	Smox emitted	Abatement cost	Smox emitted	Abatement cost
A	20	0	10	100	10	100	20	0
B	20	0	10	100	10	50	10	50
C	20	0	10	100	10	50	10	50
D	20	0	10	100	10	20	0	45
Total	80	0	40	400	40	220	40	145

otherwise A would have emitted more than 20 in the unregulated case (for free). This observation suggests an upper-bound to the bunching that may occur under emissions trading: it is highly unlikely that sources will purchase allowances to raise their emissions above their unregulated levels. Given that purchasing allowances (or paying the tax) is costly, it seems almost inevitable that bunching cannot exceed the unregulated (free) level of emissions, all other factors held constant. Of course, a source's unregulated business-as-usual emissions might rise over time as its product output grows, so a source that initially emitted 20 might later want to emit 25 or more. Adopting a fixed emissions limit would prevent that growth, but technology mandates could permit such growth to occur at all sources, and tradable allowances could permit such growth to occur at some sources while reductions (to meet the aggregate cap) occur at other sources. Also, if emissions trading were introduced as a replacement for previously imposed fixed emissions limits, then emissions at source A might grow to exceed the previous (regulated) level, but emissions at other sources would have to fall more sharply as A purchased allowances. Usually though, emissions trading is adopted to meet a more stringent cap that reduces aggregate emissions considerably below previous requirements, so increases at any source above the previous level are unlikely.

So far the analysis has ignored the dose–response function; it has implicitly assumed that every unit of exposure is equally harmful (constant marginal harm from exposure), and that every unit of emissions yields one unit of exposure (constant marginal exposure from emissions). In that case, aggregate harm depends only on aggregate emissions.

The claim is often made, however, that under the locational flexibility allowed by emissions trading, ‘bunching’ of higher emissions in one or some locations, like source A in the example above, could create ‘hotspots’ (near source A) that increase total harm (or decrease total harm less than intended). Of course, as emissions allowances are bunched at A, they must be ‘drained’ from other locations where emissions are overcontrolled to sell allowances, like source D in the example above, and perhaps the net effect of draining and bunching is a wash. (The same pattern of bunching and draining could occur under emissions taxes.)

But the situation is more complex and depends on the dose–response function illustrated in Figure 1. Consider four possibilities.

(1) Under a linear dose–response function (each unit of exposure causes one unit of harm), with constant marginal harms of exposure, and constant

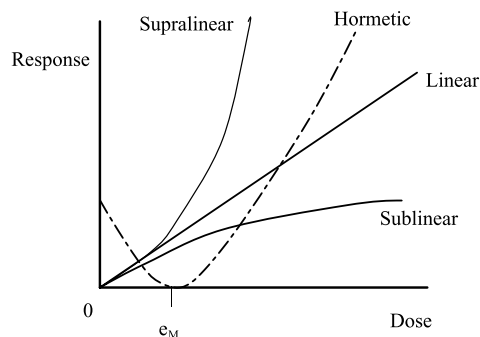


Figure 1 Dose–response curves

marginal exposure from emissions, bunching and draining (heterogeneity in emissions) would not increase total harm compared to a uniform distribution. They would offset each other. This is the standard initial case analysed above.

(2) Under a supralinear dose–response function, which is monotonically increasing (not hormetic) but has increasing marginal harms (each added unit of exposure causes more harm than the prior unit), and with constant marginal exposure from emissions, bunching would yield rising harms, but draining would yield declining harms, compared to a uniform distribution of emissions. Thus in the example above, under emissions trading as compared to the fixed emissions limits of 10 at each source, the increase at source A from 10 to 20 would raise harms more than if A and D each emitted 10; but the decline at source D from 10 to zero would also reduce harms more than if A and D each emitted 10. The net effect is uncertain, and depends on whether the buyers (bunching) and sellers (draining) are above or below the average level of emissions. (a) If buyers increase emissions above the average and sellers reduce emissions below the average, and the dose–response function has increasing marginal harms, then the harms from bunching above the average will likely exceed (grow more steeply than) the benefits from draining below the average. Divergence of buyers and sellers away from the average of the supralinear dose–response function will increase harms on net. (b) If however buyers increase emissions from below the average toward the average, and sellers reduce emissions from above the average toward the average, then the result will be convergence toward the average and the benefits of draining may exceed (harms fall more steeply than) the harms from bunching. Convergence of buyers and sellers toward the average of the supralinear dose–response function will decrease harms on net. (In the smox example above we began with uniform emissions of 20 at each source, but in reality the buyers and

sellers could start with different emissions, and then either diverge or converge through emissions trading depending on their costs of abatement.)

The supralinear dose–response case might describe the SO₂ emissions trading system adopted in 1990 to reduce acid rain. Some sources of SO₂ in the midwest emit plumes that yield acid deposition in relatively sensitive ecosystems such as the Adirondack Mountains, while other sources in other parts of the USA emit plumes that land in relatively insensitive areas such as the Atlantic Ocean. If the emissions trading programme had resulted in undercontrol (bunching) by allowance buyers whose emissions are deposited into the sensitive Adirondacks (above average, steeper marginal harms), and overcontrol (draining) by allowance sellers whose emissions would thereby be removed from eventual deposition into the insensitive Atlantic (below average, flatter marginal harms), the net result could have been an increase in harm compared to fixed emissions limits. Congress was aware of this concern but nonetheless adopted a national SO₂ emissions trading market for several reasons, including the prediction that the pattern of actual control costs at SO₂ sources would yield the reverse: more allowance selling (draining) at sources that deposit into the Adirondacks and more allowance buying (bunching) at sources that deposit into the Atlantic or other places – a happy coincidence of control costs and wind patterns yielding a net benefit from emissions trading. That is, the SO₂ market resembled case 2(b) above, with draining occurring on a steeper part of the dose–response function, and bunching occurring on a flatter part of the dose–response function, hence yielding net benefits despite any bunching. Another consideration favouring the adoption of the SO₂ trading programme was the prediction that the aggregate 50% reduction in emissions from the 1987 level would help the Adirondacks far more than any bunching might detract from that benefit.

The supralinear dose–response case could also describe an emissions trading programme in which some emissions affect sensitive subpopulations who are increasingly damaged by higher emissions. Some have raised this point as an environmental justice concern about sensitive subpopulations – including low-income and minority populations – who are residents located near sources of toxic materials, such as emissions trading involving mercury or industrial chemicals.⁵ The same two factors that made trading a net benefit in the SO₂ case could also be applicable in the environmental justice cases, but that is an empirical question, depending on the pattern of control costs and thus of buying

(bunching) and selling (draining) and the corresponding marginal harms.

The basic point here is that the concern about the harmfulness of bunching and hotspots must arise, if at all, from a supralinear dose–response function. If the dose–response function is linear, as noted above under (A), then bunching and draining due to emissions trading would have no net effect on harmfulness. (Fairness is another matter, to which I return below.)

(3) If the dose–response function is sublinear, it flattens out and has declining marginal harms with increasing exposure. In this case bunching (from buying allowances) and draining (from selling allowances) would have the converse orientation to the supralinear case. In the sublinear case, trading would be beneficial on net if buying (bunching) and selling (draining) diverge from the average, so that bunching occurs where the dose–response curve is rising but flattening out, and draining occurs where the dose–response curve is falling more steeply. But trading would be harmful on net if buyers and sellers converge toward the average of the sublinear dose–response curve, i.e., if bunching occurs where the dose–response function is rising more steeply and draining occurs where the dose–response curve is falling slowly.

(4) Under a hormetic or J-shaped dose–response function, harm increases with increasing exposure above a certain positive exposure level (the nadir of the function, labelled e_M in Hammitt's graphs for the harm-minimizing exposure³), but harm also increases with decreasing exposure below that point (e_M). If there are constant marginal exposures from emissions, and if the regulatory standard (the number of allowances) is set precisely at e_M , and if all sources begin at e_M , then under hormesis, bunching (undercontrol to emit above e_M via buying allowances) would yield increasing harms *and* draining (overcontrol to emit below e_M and selling allowances), would also yield increasing harm compared to fixed uniform emissions by all sources at e_M . Under these restrictive conditions, bunching is harmful because it ascends the positively sloped portion of the J-shaped curve to the right of e_M ; this is the same problem as under a supralinear but monotonically increasing dose–response function (when bunching moves sources above the average, as discussed above under case 2(a)). Under the above restrictive conditions, the twist is that draining (overcontrol) is *also* harmful because it ascends the negatively sloped portion of the J-shaped curve to the left of e_M .

If, however, emissions trading starts with some sources emitting above e_M and some emitting below

e_M , and if trading results in buyers bunching from below e_M upward toward e_M , and sellers draining from above e_M downward toward e_M (which depends on their starting points and their costs of abatement), then emissions trading would converge to e_M and to the level of harm under fixed emissions limits set at e_M .

In our smox example above, if the harm-minimizing level of exposure (e_M) occurred with emissions of 10 at each source (note that this assumes local effects of emissions), then the unregulated emissions of 20 at each source would be harmful, and the move to fixed emissions limits of 10 at each source would be optimal. Under emissions trading, the increase at A from 10 to 20 would be harmful, and the decline at D from 10 to zero would *also* be harmful because under hormesis the harms also rise as exposures fall below e_M . If 10 is the optimal harm-minimizing level of emissions (e_M) (and emissions translate directly into exposures), and if each deviation of 10 emissions units in either direction increases harm by, say, 10 units of harm (a V-shaped hormetic curve which is linear at 45° angles on each side of e_M), then before regulation the aggregate harm from emissions of 20 at each of four sources was $4 \times 10 = 40$. With fixed emissions levels of 10 at each source, the aggregate harm is zero. With emissions trading as illustrated in the smox example above, the harm is 10 from A emitting 20, zero from each of B and C emitting 10, and 10 from D emitting zero, for a total harm of 20. If the hormetic curve is flatter than this (rises more slowly) on each side of e_M , the aggregate harm under emissions trading would be smaller than 20.

Hormesis and the local effects of emissions are distinct attributes

The first implication of this analysis is that the effects of hormesis must be distinguished from the local effects of emissions. Some pollutants may be hormetic and have local emissions effects, while others may be hormetic without their emissions causing local effects. Hormesis is an attribute of the dose–response or exposure–response relationship. Hotspots are an attribute of the emissions–exposure relationship. It is only the local impact of emissions that poses a problem for emissions trading. In the example above under case (4), emissions trading yielded increasing harms under hormesis at both source A (increasing emissions) and source D (decreasing emissions) only because we assumed that the optimal emissions at each source (e_M) were 10 and that deviations from that point would yield harms. Some pollutants pose no such local effects

from their emissions. If smox affected the lake only because of its aggregate concentration in the entire lake and not its local concentration near each source, then the reduction from 80 to 40 in the lake would still be far above the e_M of 10 and on the positively sloped portion of the dose–response curve, and the emissions trade between A and D would not affect harms at all. Whether emissions were at 20 or 10 or 0 at each source would not matter to harms by itself; only the total smox in the lake from all sources combined would matter.

Hammitt cites CFCs and CO₂ as substances that mix globally in the atmosphere and therefore the emissions of which have no local effects: the harms of these emissions are utterly independent of where they are emitted and depend only on the total quantity in the earth's atmosphere.¹ The ultimate effects of ozone depletion and global warming would be geographically uneven and would have larger effects in some locales than others (e.g., the ozone 'hole' over Antarctica, more skin cancer in some populations than others, greater changes in temperature and precipitation in some places, and greater harms from sea level rise in some places), but these impacts would arise from global processes after the CFCs and CO₂ are emitted and mixed into the global atmosphere, and would not be affected by the geographic pattern of the initial emissions. (A caveat: extremely large local releases of CO₂ could have local effects, such as asphyxiation of ocean life due to sudden sea floor outgassing, but such releases are extremely rare and are not the kind associated with fossil fuel combustion or emissions trading.) Because their harms of emissions (and benefits of abatement) are independent of the location of emissions, both CFCs and CO₂ are therefore excellent candidates for emissions trading, including international emissions trading. The USA used a tradable allowance system to help phase out CFCs. The costs of CO₂ abatement vary considerably, so the cost savings from emissions trading are quite large.⁶

But although CO₂ emissions have no local effects to speak of, CO₂ does have a hormetic dose–response function (that is, where response means harm; CO₂ seems to have a monotonically increasing effect on temperature, but a hormetic effect on human and ecological welfare). Whatever one thinks of predictions that rising levels of CO₂ in the atmosphere may cause global warming, it is fairly clear that declining levels of CO₂ in the atmosphere below some level would also be harmful. The preindustrial level of CO₂ in the atmosphere, circa 1800, was about 275 ppm; the current level in 2003 is about 375 ppm. Exposure rising above 400 ppm is predicted to raise global average temperature. But

exposure below, say, 250 or 200 or 150 ppm could lower global average temperature dangerously and also suppress plant photosynthesis. Just as the historical record suggests that rising CO₂ is associated with warming, it also suggests that falling CO₂ is associated with cooling and ice ages. Cause and effect are not fully established – perhaps warming and cooling cause changes in CO₂ concentrations rather than the other way around – but if CO₂ influences the climate, then it clearly exhibits hormesis, and the optimal exposure (e_M) is probably somewhere around 250–350 or 450 ppm (although this point is controversial and difficult to test), with rising harms on either side.

Thus, hormesis alone does not pose any problems for emissions trading. A necessary condition for hormesis to detract from the success of emissions trading is that emissions themselves have local effects. The ‘where’ flexibility offered by emissions trading is of concern only when the location of emissions matters to harm. Not all cases of hormesis are also cases of emissions-dependent local effects.

Dependence on the level of protection

The above analysis shows how emissions trading can under certain restrictive conditions perform less well or even perversely if the standard is set at the nadir of the hormetic curve (e_M). In the example above, case (4), emissions trading yielded increasing harms (compared to fixed emissions limits at each source) because the harms were dependent on local emissions, and because both the buyer (A) and seller (D) deviated from the optimal emissions point (e_M) at 10. Emissions-dependent local harm, however, is only one of the necessary conditions for this perverse result to be obtained. In addition, it is also necessary that the standard has been set at or near the harm-minimizing point (e_M) on the hormetic curve.

If the standard had been set at a higher level than e_M , such as at the point Hammitt labels e_L^* or e_H^* (the optimal points for the linear and hormetic functions, respectively, when setting marginal benefits equal to marginal costs of control), then modest deviations from the standard by emissions allowance buyers (bunching) and sellers (draining) would still be on the positively sloped portion of the dose–response curve, and the analysis would be similar to that for the supralinear curve as described above in case (2). Bunching would be harmful but draining would be beneficial, and the net effect would depend on where the bunching and draining occur on the curve. If by contrast the standard was set

below e_M (such as a very stringent level of control for, say, red wine or vitamins, set below the optimal level of that substance for long-term health), then draining would be harmful but bunching would be beneficial, up to e_M .

Thus, in addition to depending on emissions having local effects, the influence of emissions trading on net benefits under hormesis depends fundamentally on what level of protection is set and how that level compares to the nadir of the curve (e_M). The perverse results are worst when the standard is set precisely at e_M , and the perversity diminishes as the standard moves away from e_M .

Yet it is unlikely that the standard would be set at e_M . As Hammitt shows, if there are positive marginal costs of abatement, then under hormesis the optimal level of control is still above (to the right of) e_M , at a point like e_H^* . On that analysis, the point e_M would only be the optimal point to set the standard if the costs of abatement were zero, which seems implausible. At e_H^* , the hormetic curve is supralinear, and thus the impacts of draining are still beneficial (so long as they do not go so far as to slip below e_M). In the example above, if e_M were at 3 units of smox emissions and the government set the standard at $e_H^* = 10$, and if A decided to buy only 6 units from D rather than all 10, then the resulting emissions would be A = 16, B = 10, C = 10, D = 4. A’s bunching would increase harm above the alternative of fixed emissions limits for each firm at 10, but D’s draining would reduce the harm compared to 10. Only if D sells more than 7 would D slip below $e_M = 3$ and incur the perverse harms of too few emissions. It would be useful to review the emissions trading programmes adopted to date, such as for leaded gasoline, SO₂ (acid rain), NO_x, and now CO₂, to see if any of these programmes set the standard so low near e_M (and also had effects dependent on local emissions) such that locational flexibility risked perverse increases in harm from both bunching and draining; and if the standards were set above e_M , then the effects of bunching and draining should still be netted out along the linear or supralinear dose–response function.

Some policies may of course set the level of protection at or below e_M , even if considering cost would suggest a less stringent standard. If the standard for a hormetic substance or agent is set below e_M , thereby depriving society of the health benefits of increased exposure up to e_M , it is unlikely that emissions trading would be used to implement the policy because the level would be so close to zero that there would be little scope for overcontrolling emissions and selling allowances. Moreover, even if emissions trading were employed

for a standard set at or below e_M , it could turn out that the sellers were the larger emitters (above e_M), and buyers were the lower emitters (below e_M), so that emissions trading yields convergence to the optimal e_M rather than divergent bunching and draining away from e_M .

Thus the only situation in which hormesis could render emissions trading ineffective or perverse is when emissions have local effects, when the standard is set at or near e_M , when e_M is sufficiently far from zero to enable selling of allowances, and when allowance sellers are emitting below e_M and buyers are emitting above e_M .

Further, the relevant comparison is not just to the alternative of fixed emissions limits at each firm, but also to the level of unregulated emissions (e_U). The aggregate reduction in harm from unregulated e_U to the regulated level would often dominate any increases in harm that might ensue from the heterogeneity in local emissions allowed by emissions trading. That is, the difference between e_U and the standard set at e_H^* (or e_L^*) may be so great that any bunching and draining around the standard would be negligible in comparison. This was one of the reasons for establishing a national SO_2 trading system in 1990, as noted above. If the benefits of the aggregate reduction from e_U to the standard are large enough, they can outweigh any harms from locational flexibility at every locale, thus yielding net benefits at every locale. Of course, fixed emissions limits could try to ensure greater net benefits at the sensitive locales subject to bunching. But fixed emissions limits would forego the benefits of draining at other locations (unless the draining slips below e_M), and they would forego the cost savings from emissions trading. Because the cost savings from emissions trading may be the key to enabling government to set more stringent standards than under fixed emissions limits (as they arguably were in the SO_2 case), the increase in protection of sensitive locations from employing fixed emissions limits may be outweighed by the increase in aggregate benefits to all locations (including the sensitive locations) from employing emissions trading. To illustrate using the SO_2 case: if the cost savings obtained by shifting from fixed emissions limits to emissions trading enabled the programme to achieve, say, a 10 million ton cut in national SO_2 emissions rather than a 6 million ton cut in emissions at the same or lower overall cost, but emissions trading risked a 0.5 million ton increase due to bunching at sources that would deposit in the Adirondacks (although actually the pattern of abatement costs suggested draining at such sources, as noted above), then the extra 4 million tons of

aggregate emission reduction from emissions trading as compared to fixed emissions limits could yield a net benefit to the Adirondacks despite the bunching. To illustrate using the smox example above: if the cost savings from emissions trading (\$75) were partially or fully used to obtain more stringent emissions reductions (i.e., by spending up to \$220 under emissions trading to equal what would have been spent under fixed emissions limits), thus reducing aggregate emissions not just to 40 but, say, to 24 (6 allowances allocated to each source instead of 10), and if smox is not hormetic and its e_M is not low, then the extra harms to the area near source A from any bunching due to A's purchase of, say, the 6 allowances from D could be outweighed by the benefits to all – and even to A alone – from the extra aggregate reduction achieved (16 units).

Fairness

Bunching and hotspots may be of concern for two possible reasons, harmfulness and fairness. So far we have addressed only harmfulness. Fairness concerns may remain even if harmfulness is not affected. In our initial smox example, fixed emissions limits left each source with 10 units of emissions, whereas emissions trading left A with 20, B and C with 10 each, and D with zero.

Assume that the dose–response function is linear (with one unit of smox generating one unit of harm), and that local emissions have local effects. Thus under no regulation, the harms are 20 at each of the four source locations and 80 in total. Under fixed emissions limits, the harms are 10 at each and 40 total. Under emissions trading, the harms are 20 at A, 10 each at B and C, and zero at D; the total harm remains 40. This is case (1) analysed above in which emissions trading makes no difference to aggregate harm; the dose–response function is not supralinear or sublinear or hormetic, so bunching and draining have no disproportionate effect on harms. Harms are 40 under both policies, and no source has increased its total emissions above the unregulated level of 20. But now the local harms differ, and the community near source A may complain of unfairness because it suffers harms of 20 while the communities near B and C suffer 10 each and the community near D suffers zero. (The community could be the human population or, say, the aquatic population of organisms living in the lake near the source's discharge pipe. Whether a nonhuman ecosystem can raise fairness concerns is another story; for now I assume it can.)

For example, perhaps smox causes impaired childhood brain development (as does lead (Pb)) and attendant lost IQ, so that before all four communities were equally (and seriously) impaired, but now A's children will gradually fall behind B, C and D over time. The community at A might feel that this is unfair even if A does not increase emissions above its preregulation level of 20 (which as argued above it would not), so that the community near A is at least experiencing no greater emissions than it did before the policy. Although the other communities now experience lower emissions than they did before the policy, they also incur the costs of abatement. And the aggregate benefit is significant. Still, equality of future opportunity (freedom from unequal smox impairment) may be worth something to each community, in addition to the aggregate reduction in smox. If so, although the community near A is not experiencing increased emissions, it is harder to say that it is no worse off than before, given that it has not enjoyed any reduction in emissions (it is still at 20) and now it will lag behind the other three communities in IQ over time (and its source is paying to purchase some allowances – but that might save jobs in A that are being cut in D). This fairness concern might be restated as a claim that the policy is not Pareto-improving for each community. On the other hand, it is hard to argue that every community should be exposed to equal risk. No-one argues that police or fire protection should aim to ensure equal mortality rates across all communities. If the cost savings from emissions trading enabled a tighter overall standard, then foregoing trading to maintain uniform fixed emissions levels would be less unfair to community A at the cost of greater injury to all (more impairment of other children). Perhaps community A could be compensated through side payments or through the income tax system. People near A could use the exit option to move to B, C or D. Or community A could use its voice to induce source A to reduce emissions further, or the government to tighten the aggregate level even further.

If the dose–response function is supralinear, such as because of more sensitive subpopulations at some locations, and if those subpopulations are subject to higher emissions due to bunching, then those subpopulations may raise a fairness concern that is compounded by increasing harmfulness. This would be the case if the community near A was also sensitive to greater injury from smox than other groups. This fairness concern is more compelling than the fairness concern associated with linear damages, because now the fairness concern is combined with a greater vulnerability to harm.

This is the type of fairness concern often posed in environmental justice cases, where a subpopulation with lower incomes, perhaps lower nutrition and education, and perhaps victims of historical discrimination on the basis of race or ethnicity, argues that it would be subject to a higher exposure and that it is more vulnerable to damages from each unit of exposure. Of course, such concerns can be raised wholly apart from emissions trading systems, such as the siting of new industrial facilities.

On the other hand, if the sensitive subpopulation is located where the source subject to emissions trading is not bunching but draining, such as at source D, then the combined fairness and harmfulness concern would be neutralized or reversed. This was predicted to occur for the Adirondacks under the acid rain trading programme, as lower-abatement-cost SO₂ sources in the midwest would overcontrol (drain) and sell allowances to higher-abatement-cost sources elsewhere which would undercontrol (bunch) and emit into other less sensitive deposition sites.

If the dose–response function is hormetic, and if the standard is set at or near the harm-minimizing level (e_M), then the fairness concerns are complicated by the possibility of harm from both higher and lower exposures than e_M under emissions trading. As discussed above, however, it seems unlikely that the standard would be set at e_M .

Now assume that the harms depend only on the aggregate level of smox in the lake, not on the local emissions at each source. Then the community near source A should have no fairness concern. A does emit more than the others, but the local community is not affected any more or less by these emissions than everyone else. Indeed, the community near A may be accused (rightly or not) of unfairness by others around the lake who incur the harms of the aggregate 40 and see A emitting fully half that. (Of course A did so quite legally pursuant to the government's policy design, and the aggregate reduction to 40 – which is all that matters in this case – has been achieved at cost savings to A and profits to D compared to fixed emissions limits.)

From emissions to risk

To the extent that hormesis poses a problem for market-based incentives such as emissions trading and emissions taxes, which as I have argued above is limited, the solution to these problems may not be to abandon market-based incentive instruments and their cost-effectiveness gains, but to improve them further by moving from emissions trading and emission taxes to risk trading and risk taxes.

One way to avoid the potentially perverse effects of emissions trading or taxes under hormesis is to forego emissions trading and taxes, and stick with fixed emissions limits. As noted above, that choice forfeits the cost savings and more stringent standards made possible by market-based incentives, which may well be more important than the possible perversities around the hormetic optimum. A second option is to constrain the market by limiting bunching (acquiring extra allowances) and draining (selling allowances), such as by forbidding buyers to buy more than a certain amount, or sellers to sell more than a certain amount, or anyone to emit more than some maximum and minimum. Those options may be helpful but also inhibit the market and curtail its cost savings. They are partial moves toward fixed emissions limits. A third option is to segment the market geographically so as to prevent trades that cause bunching and draining in certain places. For example, if the community near source A is especially vulnerable, the policy could allow emissions trading but prohibit sales to A. Or the acid rain trading programme could restrict sales to sources whose emissions would be deposited in the sensitive Adirondacks (as New York State tried to do, by forbidding LILCo from selling allowances to upwind sources). These approaches also inhibit transactions and forego cost savings. Thus in each of these cases, the gains from avoiding hormetic perversities must be weighed against the costs of using the less cost-effective design.²⁷

The fundamental reason that emissions trading and taxes confront problems under hormesis is not, however, a problem with trading or incentives. It is the focus on emissions as the target of the policy. Hormesis is vexing because it reveals that emissions and exposure are not proportionate to risk. And as Hammitt notes, even without hormesis, if emissions are not always equally proportionate to exposure – if emissions translate into different exposures at different sources – then emissions trading and taxes will not outperform command-and-control policies as well as they do when emissions yield equal exposures everywhere.

The lesson is that emissions are only an intermediate stage in the production of risk. Policy can operate at several different control points, including the inputs to the firm (e.g., fuels and materials), the firm's internal processes and technologies (e.g., combustion methods, scrubbers), the firm's outputs such as emissions (e.g., emissions limits, allowances or taxes), the ambient levels of the pollutant (e.g., workplace ambient standards or state-level NAAQS or CO₂ levels), exposure (e.g., worker protective equipment) and ultimately risk (e.g., tort liability

or workers' compensation). In practice we can see examples of all of these approaches. Much air and water pollution policy has focused on the firm's technology, emissions or ambient levels. And tradable allowance policies have often focused on emissions (although the lead phasedown involved tradable gasoline content credits applied upstream to refiners – an arrangement that helped avoid any local emissions effects).

But if emissions are not closely related to exposure and in turn to risk, or are sometimes inversely related to risk because of hormesis, then policy should turn to controlling risk rather than emissions. As Hammitt notes briefly, 'A more complicated system, in which the tax or quantity of permits required per unit of exposure varies across subpopulations in proportion to the marginal benefits of reducing exposure could provide superior outcomes but would be more complicated to develop.'¹

One approach would be a tax-subsidy regime which taxes emissions (or exposure) above e_M , and subsidizes them below e_M .⁷ A second approach is to trade risk allowances, not emissions allowances.⁸ 'Risk bubbles' have been advocated within firms and across site-specific projects. For example, Paul Portney has suggested that 'regulated entities, public or private, should be allowed to relax pollution controls at one point and install new protective measures elsewhere, subject to a demonstration that an overall improvement in health or environmental quality will take place as a result of the change.'⁸ In Richard Stewart's formulation, 'Under risk bubbles, a facility operates under an aggregate, multi-media 'umbrella' of residuals limitations for the facility as a whole. This 'risk cap' creates economic incentives because higher levels of discharges by the facility of a given residual in a given medium carry an opportunity cost in the form of the resources that must be devoted to reducing other residuals in order to stay within the cap. In effect, the cap creates an implicit internal residuals trading market.'⁹ Thus a plant could relax air pollution controls at one place in exchange for achieving an extra reduction in water pollution emissions elsewhere in the plant or at another plant in the area, subject to EPA approval of the net reduction in risk. EPA and the plant owners would negotiate the bubbling arrangement,¹⁰ and the plant owners would bear the burden of proving that their water pollution reductions are at least equal to the air pollution increase in terms of environmental risk. EPA has begun to use risk bubbles to allow net risk improvements at individual firms through Project XL.^{9,11} Others have advocated letting parties responsible for Superfund cleanups negotiate agreements with host municipa-

lities and states in which all would agree to spend some of the funds on other measures that would reduce risks even more, such as air pollution control or city street lighting or vaccines for children.¹² Frank Cross advocates setting an aggregate risk limit (called a 'risk cup') and then requiring new sources of additional risk to purchase offsetting reductions in risk from existing sources of risk.¹³ The Clean Air Act's provisions for air toxics offsets (section 112(g)) and CFC substitutes (section 612) both authorize interpollutant trading based on achieving reductions in overall risk.

The concept of risk bubbles, like emissions bubbles, represents an initial project-based version of locational flexibility. And just as the experience with emissions bubbles under EPA's bubble, netting and offsets policies in the late 1970s and early 1980s became the basis for the more formal and more successful emissions allowance trading policies for lead, CFCs and SO₂ in the 1980s and 1990s, so too can risk bubbles build the groundwork for formal risk allowance trading. Under formal risk allowance trading, each source would be allocated allowances to cause a limited number of units of harm, significantly reducing harm from the unregulated level. Then each source, perhaps assisted by the government, would calculate the effects of its activities on risk and apply its allowances accordingly. Low-cost risk abaters would reduce risks further and sell allowances; high-cost risk abaters would buy allowances.

Risk bubbles and risk allowance trading would overcome the problem of bunching, hotspots and hormetic perversity – and perhaps even unfairness – because they would measure results in risk rather than in emissions. Thus if bunching emissions at location A would increase risk, such as to a sensitive subpopulation, that would require additional risk allowances, and sources would have an incentive to reduce those risks. If reducing emissions below the hormetic nadir e_M would increase risk, that would require additional risk allowances (not fewer emissions allowances). To be sure, a firm could purchase additional risk allowances, increasing the risk it poses. But as noted above, it is unlikely or inconceivable that a firm would purchase costly allowances to emit or cause risk at levels higher than it would have done before being regulated. Thus the creation of a system of risk allowances, even if initially set at the current level of risk in the aggregate (preventing increases in risk), is unlikely to yield significant increases at specific firms. And if the risk allowances are issued at a lower level than is currently obtained, net reductions could be anticipated in most places, with buyers increasing above the standard but not above the preregulation level.

In addition, risk bubbles and trading represent a vehicle for integrated multimedia pollution control,¹⁴ overcoming the problems of cross-media shifts of pollution and risk–risk tradeoffs across substances, activities and fragmented narrow regulatory programmes that have plagued our risk regulation regime.¹⁵ Provided all significant risks created by the firm are counted, the risk bubble or allowance system would remove any incentive to shift away from one regulated activity or substance or medium into another; the firm would be accountable for every unit of risk it creates, and would have incentives to reduce its overall risk in the most cost-effective ways possible. '[T]he facility is forced to internalize the environmental risks posed by its operations and take them into account when evaluating the regulatory feasibility and the economic profitability of changes in technologies and production. Any increase in one residual must be paid for with reductions of others.'⁹ It would also stimulate a wave of new innovation in risk-reduction opportunities previously ignored under the current fragmented regulatory system.

Risk allowance trading is easy to parody and pillory as selling the right to murder. But this epithet is misplaced. First, risk allowance trading would not authorize intentional assaults or any other forms of risk creation that are already illegal. Criminal law prohibitions would always supersede regulatory risk allowances. Secondly, risk allowance trading would be used to sharply reduce risks from existing levels, protecting society against ills not yet addressed, and at lower cost than fixed limits on risk. Thirdly, it would solve the problem of exposure to more sensitive subpopulations and the perversities of emissions trading under hormesis, thereby protecting society even better than emissions limits, taxes and trading systems do. Fourthly, it would reinforce and improve on the tort law system, which provides *ex post* financial liability remedies for risk, by adding an *ex ante* quantity limit on risk – a more protective and credible regime, set at levels determined by comprehensive analysis of societal risks, without the high transaction costs, variabilities and unpredictabilities of litigation.

At the same time, risk bubbles and risk allowance trading would have to satisfy important criteria. They demand information about risk levels imposed by different activities, and some means of making those risks commensurable across different substances, activities, populations and health endpoints. Stewart notes: 'Each such facility is given an individual risk quota; a maximum level of risk that it may impose on the environment as a result of the residuals that it generates. The foundation of the 'risk bubble' system is

a set of risk indices assigned to each residual and measured in some common risk unit. A common risk unit may, for example, be fixed as a level of the environmental hazard created by a discharge of one pound of a reference pollutant, the effects of which on human health and the environment are well known. The relative risk measure of other regulated pollutants are determined and assigned in relation to the baseline pollutant. An individual facility is free to select the levels of specific residuals that it generates as long as the aggregate risk imposed on human health and the environment by all of the total discharge does not exceed the facility's maximum risk quota.⁹

This will not be easy. The creation of 'a comprehensive system for indexing the relative environmental risks of different residuals... presents a problem of serious scientific difficulty.⁹ We would need better information on emissions, associated exposures and associated risks for a host of substances and activities, including their synergistic or offsetting effects in concert. Mechanisms, dose-response functions, exposures and variability across populations would need to be better understood and measured. Many uncertainties in current science would need to be ameliorated. Diverse risks would need to be translated into some version of commensurability – some set of multiattribute risk metrics that can enable transparent accounting and review. And improved monitoring, perhaps by third parties as well as by government, would be needed to ensure that firms' risks are accurately counted, reported and checked against their allowances.

Yet these challenges are not unique to the risk bubble/trading approach; they apply equally to the current regime of fragmented command-and-control and emission trading systems. Not indexing is tantamount to indexing arbitrarily. 'The relative stringency of command requirements for control of different residuals discharged into different media represents a de facto index of relative risk and regulatory priority. The virtue of a risk bubble approach is that it makes the indexing issue explicit and focuses scientific research and policy work on improving such values and dealing explicitly and systematically with risk versus risk and benefit versus benefit tradeoffs.⁹ Moreover, commensurability across risks is an endogenous function of our investment in measuring and understanding, rather than an inescapable extrinsic barrier.^a The familiar

complaint that someone is inappropriately 'comparing apples and oranges' actually disproves itself: people regularly do compare apples and oranges when choosing what to buy at the grocery store; the statement that apples are different from oranges itself constructs the comparison on which a choice could be made (such as based on flavour, colour, nutrient content, price and other attributes); and comparing the difference between 'apples and oranges' to the difference between the other two things that the speaker is arguing cannot be compared (such as two risks) is itself a comparison that proves its feasibility.¹⁶ Indices of the relative impacts of different pollutants can be constructed to enable transparent comparisons across pollutants contributing to the same risk and, potentially, different risks contributing to overall harm.^b

Likewise, inadequate monitoring impairs technology-based and emissions-based regimes as well as risk-based regimes. Who knows if the scrubber is turned on, or operating effectively? Investments in better monitoring are needed in all regulatory systems. The risk bubble/trading approach would direct incentives toward (private sector) investment in better monitoring of the activities and attributes that actually contribute to harm. Dan Esty argues that although risk bubbles 'depend on a greater degree of information than is generally currently available... Today, however, with advanced pollution detection and tracking equipment, the technical dimension of the risk bubble problem is becoming increasingly tractable.'²² Likewise, Donald Elliott and Gail Charnley stress the practicality of designing effective risk bubbles and trading systems, if investments are made in the science of monitoring, measurement and comparison.^c Regulatory regimes can be designed to reward industries' investments in

^bIndices to compare multiple pollutants contributing to a single risk have been successfully developed and used in the treaties addressing CFCs (ozone depletion potential or ODP) and GHGs (global warming potential or GWP). Indices to compare different risks with each other can draw on willingness to pay/accept approaches (using either revealed or stated preferences). Short of full monetization, such indices can be constructed using multiattribute decision theory, by combining different factors into an index through additive weighting.^{17–21}

^c'On the one hand, the more broadly the measures of equivalent compliance are defined, the greater the opportunities for flexibility, improved performance, and cost savings. Thus, a comprehensive risk bubble that would permit one kind of environmental risk to be traded off against another would potentially provide the greatest benefit. On the other hand, the difficulties of measuring 'equivalency' become greater the more broadly the concept is defined. Thus, it is inherently easier to equate one pound of sulphur dioxide released in one part of a factory with an equivalent pound of the identical chemical released somewhere else in the same facility. However, as common metrics are developed for comparison purposes, dissimilar environmental risks will be able to be traded off more broadly against one another.'^{23,24}

^a'For risks that seem very 'dissimilar', risks cannot (yet) be measured on a unidimensional scale, and the exercise of informed value judgments becomes all the more central. ... But it is chiefly our lack of methods of comparison ... that makes these risks seem 'dissimilar' or noncomparable, not an inherent incommensurability. As we improve methods of risk analysis, the idea of calculating the 'net risk' of a risk portfolio, or the change in net risk due to a risk tradeoff, may become more meaningful.'¹⁵

improved monitoring. All of these challenges pose the cost of new information against the value of new information. If health and environmental risks are important to our society, then investing in improved methods of monitoring, measurement and comparison can yield large net gains.

The legal regime for such a risk-based approach, especially risk allowance trading across firms, has yet to be designed. Project XL was unable to stimulate widespread risk bubbling within firms largely because of current legal requirements for specific controls that such risk bubbles would have had to overcome; even if EPA had agreed to waive those requirements (in return for superior risk reductions), citizens' groups could sue to enforce them. Thus firms were hesitant to agree to risk reductions in some areas when EPA could not promise relief from other rules that did little to reduce risk. New legislation would be needed to authorize full risk bubbles with intrafirm flexibility. And risk allowance trading across firms would require major new legislation, likely with special attention to whether risk allowance trading would prevent or exacerbate hotspots. A potentially greater obstacle is the motivation of interest groups and members of Congress to maintain bureaucratic control and claim credit for separate initiatives by keeping the environmental regulatory system fragmented into different laws and committees; integration may make social sense but not political sense. Gradual experience with risk bubbling could be gained through initiatives like Project XL, through practice under Clean Air sections 112(g) and 612 noted above, and through multimedia permitting. The EPA could begin developing the transparent indices and monitoring systems. Industry could be encouraged to finance this work by offering enhanced abatement credit in return for better monitoring of emissions, exposure and risk.²⁵ In addition, there is a move afoot in the current Congress to pass new 'three pollutant' air quality legislation, covering SO_x, NO_x and mercury, with a possibility of adding CO₂ as well. This legislation would authorize some interpollutant trading (e.g., SO_x-NO_x), which would build experience with the

monitoring, risk assessment and commensurability needed for risk allowance trading.

Conclusion

Hammitt is right to raise the possible problems posed for emissions trading or taxes by hormesis, but as this comment has illustrated, the problem is not as serious as might be thought. For emissions trading or taxes to be undermined by hormesis, the necessary conditions involve not just hormesis but also local emissions effects, a level of protection set at or near the hormetic minimum-effects level, and a pattern of selling and buying by sources along the dose-response curve in a direction that poses a net increase in harm. Fairness concerns about hotspots are distinct from hormetic effects. And to the extent that hormesis does pose problems for emissions trading and taxes, an eventual shift toward risk trading and taxes is the superior route.

At the same time, several issues of interest remain unaddressed. One question is how a hormetic dose-response function and a more complex emissions-exposure function would interact. Another is what an empirical review of the actual performance of fixed emissions limits and tradable emissions allowances in the past, illuminated by evidence about the dose-response function, would reveal about whether hormesis and emissions trading have conflicted in practice. A third is whether hormesis would have implications for the choice between taxes (prices) versus allowance trading (quantities) – which it might, given the importance of differently sloped marginal benefit functions for this choice.²⁶ A fourth is whether these issues bear on different methods of allocating allowances to sources. Further analysis and empirical study of dose-response functions and policy instruments could help improve the selection of instruments to match the complexities of real environmental problems.

Acknowledgement

I thank James Hammitt for helpful comments on an earlier draft.

References

- 1 Hammitt JK. Economic implications of hormesis. *Human Exp Toxicol* 2004; **23**: 267–78.
- 2 Calabrese EJ, Baldwin L. U-shaped dose-responses in biology, toxicology, and public health. *Annu Rev Public Health* 2001; **22**: 15–33.
- 3 Cross FB. Legal implications of hormesis. *Human Exp Toxicol* 2001; **20**: 122–28 (and related comments in this symposium issue).
- 4 Wiener JB. Global environmental regulation: instrument choice in legal context. *Yale Law J* 1999; **108**: 677–800.

- 5 Johnson SM. Economics v. equity: do market-based environmental reforms exacerbate environmental injustice? *Wash Lee L Rev* 1999; **56**: 111, 125–26.
- 6 Stewart RB, Wiener JB. *Reconstructing climate policy*. Washington, DC: AEI Press, 2003.
- 7 Roberts MJ, Spence M. Effluent charges and licenses under uncertainty. *J Public Econ* 1976; **5**: 193–208.
- 8 Portney PR. Reforming environmental regulation: three modest proposals. *Iss Sci Technol* 1988; Winter: 74–81, and Reforming environmental regulation: three modest proposals. *Columbia J Environ Law* 1988; **13**: 201, 207–209.
- 9 Stewart RB. A new generation of environmental regulation? *Capital Univ Law Rev* 2001; **29**: 21, 64–66, 96 (See also id. at 116–122 (analysing risk bubbles in detail)).
- 10 Pedersen WF. Contracting with the regulated for better regulations. *Admin Law Rev* 2001; **53**: 1067 (advocating government-industry contracts that apply risk bubbles to improve risk protection while lowering costs).
- 11 Farber DA. Rethinking regulatory reform after *American Trucking*. *Pace Law Rev* 2002; **23**: 43, 70.
- 12 Harvard Group on Risk Management Reform. Reforming risk regulation: achieving more protection at less cost. *Hum Ecol Risk Assessment* 1995; **1**: 183–206.
- 13 Cross FB. Incorporating hormesis in risk regulation. *Environ Law Rep* 2000; **30**: 10778.
- 14 Guruswamy L. Integrated pollution control: the expanding matrix. *Environ Law* 1992; **22**: 77–90.
- 15 Graham JD, Wiener JB. *Risk vs. risk: tradeoffs in protecting health and the environment*. Cambridge, MA: Harvard University Press, 1995.
- 16 Chodosh HE. Comparing comparisons: in search of methodology. *Iowa Law Rev* 1999; **84**: 1025, 1061–63.
- 17 Viscusi WK, Magat W, Huber J. Pricing environmental health risks: survey assessments of risk–risk and risk–dollar tradeoffs for chronic bronchitis. *J Environ Econ Manag* 1991; **21**: 32–51.
- 18 Baron J. *Thinking and deciding*, 3rd edition. Cambridge, UK and New York: Cambridge University Press, 2000.
- 19 Keeney RL, Raiffa H. *Decisions with multiple objectives: preferences and value tradeoffs*. New York: Wiley, 1976: 84.
- 20 Payne JW, Bettman JR, Johnson EJ. *The adaptive decision maker*. New York: Cambridge University Press, 1993.
- 21 von Winterfeldt D, Edwards W. *Decision analysis and behavioral research*. New York: Cambridge University Press, 1986; 259–77.
- 22 Esty DC. Next generation environmental law: a response to Richard Stewart. *Capital Univ Law Rev* 2001; **29**: 183, 196 (citations omitted).
- 23 Elliott ED, Charnley G. Toward bigger bubbles: why interpollutant and interrisk trading are good ideas and how we get there from here. *Forum Appl Res Public Policy* 1998; Winter: 48.
- 24 Elliott ED. Toward ecological law and policy. In Chertow MR, Esty DC eds. *Thinking ecologically: the next generation of environmental policy*. New Haven: Yale University Press, 1997: 184.
- 25 Wiener JB. Solving the precautionary paradox: policy approaches to improve measurement of greenhouse gas sources and sinks. In van Ham J, Janssen LHJM, Swart RJ eds. *Non-CO₂ greenhouse gases*. Dordrecht, Netherlands: Kluwer Academic Publishers, 1994: 527–31.
- 26 Weitzman M. Prices versus quantities. *Rev Econ Stud* 1974; **41**: 477.
- 27 Nash JR, Revesz RL. Markets and geography: designing marketable permit schemes to control local and regional pollutants. *Ecology Law Quarterly* 2001; **28**: 569.