The Double Dividend and Other Applications of "Normalization" as a Benchmark for Environmental Policy

Yoram Bauman

June 4, 2004

Address after August 1, 2004:

Dept. of Economics

Whitman College

345 Boyer Ave.

Walla Walla WA 99362

Address until August 1, 2004: 1225 NE 61st St.; Seattle WA 98115

Email (always good): yoram@smallparty.org

The Double Dividend and Other Applications of "Normalization" as a Benchmark for Environmental Policy

Abstract

Although the ultimate goal of environmental policy is to improve social welfare, a useful benchmark is the "normalization" of environmental goods, i.e., policies that result in parity between environmental goods and "normal" goods. This paper applies the normalization benchmark to the areas of the double dividend hypothesis, asymmetric information, dynamic optimality, and regulation of monopoly. One result is a more appropriate description of the contentious double dividend hypothesis. This description, which supports the original claims of double dividend proponents, compares a world with an environmental good with an otherwise identical world containing a similar but "normal" good.

The distinguishing feature of environmental goods is the divergence of private marginal costs from social marginal costs. For any environmental good e, then, one can imagine a sister good n that has the same social marginal cost structure as e but a "normal" private marginal cost structure, i.e., one in which private marginal cost equals social marginal cost.

Thought experiments involving e and n provide insight into a number of issues in environmental economics, including the double dividend hypothesis, asymmetric information, and dynamic optimality. In addition to shedding light on each of these three issues, this paper also hopes to make a larger point about the value of normalization—i.e., inducing firms and other relevant agents to treat e as they would treat n—as a benchmark for environmental policy.

1 The Double Dividend

Imagine sister worlds E and N, identical except that the "environmental" world E contains the environmental good e and the "normal" world N contains the normal good n.

Given that e and n have the same social marginal cost structure and that the worlds are otherwise identical, it is obvious that both worlds can achieve the same level of social welfare under ideal circumstances. Formally, this yields

Proposition 1 *The first-best outcome in worlds E and N are identical.*

Now imagine that circumstances are less than ideal, in that there is an (exogenously specified) need for government revenue and an inability to levy lump-sum taxes. How do the second-best outcomes in the two worlds match up?

One might think that a social planner in the environmental world would be lucky to reach the same level of social welfare as that attainable in its sister world. But what if the social planner in the environmental world can reach a *higher* level of social welfare than that achievable in the normal world? In this case one might argue that the presence of the environmental good e is (at least potentially) a blessing in disguise. One might further argue that, when combined with the appropriate (second-best) policies, the presence of the environmental good yields a "double dividend": the social planner in the environmental world can not only match the level of social welfare in the normal world (the first dividend), but actually surpass it (the second dividend).

A simple representative agent model shows how this double dividend can appear. Consider the second-best outcome in the normal world, which is achievable through some tax vector $\mathbf{t}$. The market price vector $\mathbf{p}$ can then be written as $\mathbf{p} = \mathbf{m} + \mathbf{t}$, where $\mathbf{m}$ is the price excluding the tax. In particular, for the sister good n we have $p_n = m_n + t_n$. (Note that in equilibrium m_n is the private—and the social—marginal cost of production.)

Next consider the environmental world E. For simplicity, assume that the private marginal cost of producing good e is zero. Now imagine setting the tax on good e equal to $t_e = m_n + t_n$, and the taxes on all other goods equal to those in world N. Further imagine instituting a lump sum transfer of $m_n \cdot q_n$ to the representative agent, where q_n is the amount of good n purchased in equilibrium in the normal world. This leads to

Proposition 2 *Social welfare in the environmental world can match the second-best level of social welfare in the normal world.*

Proof. The policy described above effectively eliminates the differences between the two worlds by equalizing the market price vectors and the representative agent's income in the two worlds.[1] The policy also equalizes government tax receipts in the two worlds: the additional taxes in the environmental world are exactly offset by the additional transfers. It follows that social welfare in the two worlds will be equal.

This shows that the social planner in the environmental world can reach the same level of social welfare as the social planner in the normal world. But clearly she can do even better than that: instead of providing lump-sum transfers, the social planner could reduce distortionary taxes. (This is nothing more than the "weak double dividend" hypothesis as formulated by Goulder 1995.) Combined with the previous proposition, the result is

Proposition 3 *Social welfare in the environmental world can exceed the second-best level of social welfare in the normal world.*

This result does not "prove" that the double dividend hypothesis is correct,

[1] Note that the additional tax collection of $m_n \cdot q_n$ in the environmental world is a surrogate for the private payment of that amount in the normal world, and also that the lump sum transfer of $m_n \cdot q_n$ in the environmental world is a surrogate for the private receipt of that amount in the normal world.

in part because there are numerous competing definitions of that hypothesis.[2] But this result does arguably support the original ideas put forward by double dividend proponents. Also clear is the value of the thought experiment involving the environmental world E and the normal world N: comparing these two worlds provides a benchmark against which to evaluate the effect of externalities on social welfare.

2 Asymmetric Information

Weitzman (1974) established the importance of asymmetric information in environmental policy-making. In particular, his work focuses on regulatory uncertainty about the demand curve for emissions, which can be related to the demand curve for the polluting good e.

Considering this issue anew in the context of the environmental world E and its sister world N—identical except that the former contains the environmental good e while the latter contains the normal good n—a natural question is whether similar asymmetries exist in the normal world N.

[2] The recent literature focuses on the size of optimal Pigovian tax rates relative to marginal environmental damage, e.g., Jaeger (2001), Goodstein (2003). But the connection between this issue and the possibility of "double" dividends is unclear. Goulder (1995) and others distinguish between an "environmental" dividend and a "non-environmental" dividend, with the second supposedly coming even in the absence of an externality. But as Christiansen (1996) points out, there is little support for the double-dividend hypothesis thus defined: absent an externality, no motivation exists for shifting taxes between goods.

The answer is not entirely clear. Although there are no asymmetries in a competitive equilibrium—each price-taking firm faces a perfectly elastic demand curve at the market price—the process by which markets reach equilibrium is shrouded in mystery.[3] In a disequilibrium context such as the aftermath of a change in input prices, a strong case can be made that the suppliers of good n in the normal world face at least as much uncertainty as the regulator faces in the environmental world: both are ill-informed about demand for their product.

It is therefore puzzling that the same economists who have great confidence in the workings of the normal world should be so concerned with the presence of asymmetric information in the environmental world. Can't poorly informed regulators in the environmental world accomplish as least as much as the poorly informed market in the normal world?

Indeed they can, and they can do so by mimicking the behavior of the supply side of normal markets. One plausible driving force in the normal world, for example, is consideration of marginal conditions: as long as the marginal benefit of good n is not equal to its marginal cost, the potential for mutual benefit provides a strong incentive for trade. Adapting this principle in the environmental world indicates that regulators can address asymmetric information problems by focusing on policies that provide firms with opportunities to pollute more (or less) if their marginal benefit of emissions exceeds (or falls short of) marginal environmental damage.

Three policies that accomplish this task are described in Roberts and Spence

[3] See Arrow (1959) for one approach.

(1976), Collinge and Oates (1982), and Henry (1989, pp. 34–40). The underlying idea in all of them is that the marginal condition will be satisfied as long as profit-maximizing firms can trade an additional unit of emissions at a price equal to its marginal environmental damage.

Such a prescription is easy to incorporate into auctioned or grandfathered tradable permit systems by requiring the regulator to buy or sell an additional permit at a price equal to the marginal environmental damage of that additional unit of emissions.[4] The traditional auctioned permit system, for example, requires only two modifications. First, instead of auctioning a fixed number of permits there should be infinitely many permits. Second, the otherwise identical permits should be numbered, with permit i requiring a minimum bid equal to the marginal environmental damage of the ith unit of emissions. If n permits are purchased, a rational-expectations solution is for all permits to sell for the same price p, where p is the marginal environmental damage of the nth unit of emissions. The number of permits actually purchased is determined endogenously, and will be equal to the number of permits issued in the traditional full-information context.

[4]Difficulties arise with Pigovian tax policies because it is not possible to determine which unit of emissions is the marginal unit.

3 Optimal Incentives for Innovation

The literature on incentives for innovation in pollution control has mostly focused on the question of *maximal* incentives for innovation. Those papers that have examined *optimal* incentives for innovation have done so by comparing the private gains from innovation with the social gains. If the private gains are greater than (or less than) the social gains, the policy under consideration is deemed to provide too strong (or too weak) of an incentive for innovation.

For a graphical example, consider Figure 1's replication of the analysis of

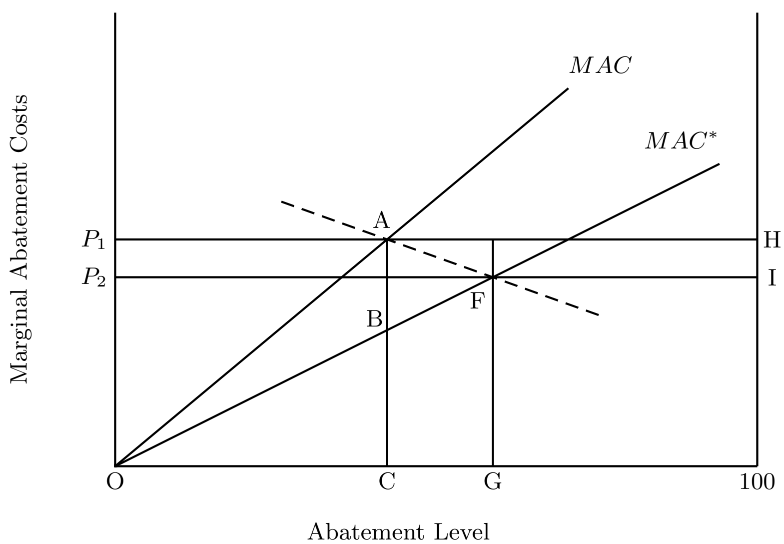

Figure 1: An innovation in end-of-pipe abatement technology which lowers abatement costs from MAC to MAC^*. The dotted line shows the marginal benefit of abatement, or—equivalently—marginal environmental damage.

end-of-pipe abatement technologies in Downing and White (1986). The innovation lowers the firm's marginal abatement costs from MAC to MAC^*, producing a shift along the marginal environmental damage curve; the regulator responds by lowering the Pigovian tax rate from P_1 to P_2. The social gain from this innovation is the area OAF. The private gain is $OAHIF$, leading Downing and White to conclude that this policy provides too strong of an incentive for innovation. The intuition here is exactly the same as in the case of a monopsonist which purchases less of an input in order to reduce the market price: by innovating, the firm can lower the market price of pollution, thereby reaping a private gain in excess of the social gain.

Instead of comparing private and social gains, the normalization approach compares the situation in the environmental world E with a similar situation in its sister world N. It is clear that the normal good n faces the same problems as the environmental good e: social incentives for innovation are higher than private incentives because innovation would lower the market price for good n. This suggests that the underlying problem in the environmental world is not the presence of an externality but the presence of market power.

This observation brings up a host of second-best questions. One issue is the relationship between static and dynamic considerations. Because the firm is a monopsonist in the market for pollution, it is likely to purchase less than the socially optimal amount. There might therefore be important interactions between dynamic and static aspects; for example, creating extra incentives for pollution-reducing innovations might alleviate (or perhaps exacerbate) resource misallocation in the static context.

Another question is whether the firm's market power is specific to the environmental good, or if it extends more generally to other input and/or output markets. In the latter case, equating private and social gains for the environmental good might be unwise because of deviations from that goal in other goods. As an extreme example, imagine that the firm has at its disposal both the environmental good e and—a perfect substitute—the sister good n. Since the two goods are otherwise identical, a firm with market power over e will also have market power over n; the divergence between social and private incentives that Downing and White identify will therefore affects both goods. As such, aligning private and social incentives for the environmental good will produce an undesirable imbalance, with incentives for n-innovation being stronger (potentially much stronger) than incentives for e-innovation.

More generally, a significant difficulty with the social-versus-private approach is that it conflates problems specific to environmental goods with more general problems that affect many or all goods. For example, patent races or free-riding are likely to lead to a divergence between social gains and private gains. Since these are economy-wide issues, it makes little sense to use environmental policy to address them for individual goods. The normalization benchmark avoids these problems by comparing private gains from different kinds of innovation; its goal is to have environmental goods treated in the same way as other goods. It would argue that the Pigovian policy described by Downing and White *is* appropriate: it produces incentives in the environmental world that are equivalent to those in its sister world.

References

Arrow, Kenneth J. 1959. "Toward a Theory of Price Adjustment." In Moses Abramovitz et al., *The Allocation of Economic Resources*. Standford Univ. Press.

Christiansen, Vidar. 1996. "Green Taxes: A Note on the Double Dividend and the Optimum Tax Rate." Center for Economic Studies working paper 107.

Collinge, Robert A., and Wallace E. Oates. 1982. "Efficiency in Pollution Control in the Short and Long Runs: A System of Rental Emissions Permits." *Canadian Journal of Economics* 15:347–354.

Downing, Paul B., and Lawrence J. White. 1986. "Innovation in Pollution Control." *Journal of Environmental Economics and Management* 13:18–29.

Goodstein, Eban S. 2003. "The Death of the Pigovian Tax? Policy Implications from the Double-Dividend Debate." *Land Economics* 79:402–414.

Goulder, Lawrence H. 1995. "Environmental Taxation and the Double Dividend: A Reader's Guide." *International Tax and Public Finance* 2:157–183.

Henry, Claude. 1989. *Microeconomics for Public Policy: Helping the Invisible Hand*. Oxford: Clarendon Press.

Jaeger, William. 2001. "Double Dividend Reconsidered." *AERE Newsletter* 21:1–20.

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

Weitzman, Martin L.˙ 1974. "Prices vs. Quantities." *Review of Economic Studies* 41:477–491.