THE
ECONOMIC
JOURNAL

The Economic Journal, **127** (*November*), 2478–2503. Doi: 10.1111/ecoj.12396 © 2016 Royal Economic Society. Published by John Wiley & Sons, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

TAXING EXTERNALITIES UNDER FINANCING CONSTRAINTS*

Florian Hoffmann, Roman Inderst and Ulf Moslener

We consider a production economy with externalities, which can be reduced by additional firm-level expenditures. This requires firms to raise additional outside financing, leading to deadweight loss due to a standard agency problem *vis-à-vis* investors. Policy is constrained as firms are privately informed about marginal abatement costs. The optimal tax on externalities is non-linear, thus, not implementable through tradable pollution rights alone, and lower than the Pigouvian tax for two reasons: first, higher outside financing creates additional deadweight loss; second, tax-induced reallocation of resources reduces average productive efficiency. Combining taxes with grants tied to loans improves resource allocation and, thus, welfare.

We consider an economy where financially constrained firms must invest to reduce externalities from production. The amount of external financing that firms raise interferes with productive efficiency due to a standard moral hazard problem. We show how in such an environment, the optimal linear tax on externalities – or, likewise, the optimal amount of tradable pollution rights – differs from the Pigouvian tax and how there is scope for additional policies, such as loan-based grants. As we discuss below, such grants are frequently observed in practice, in particular related to investments into CO_2 reduction, which should indeed be sizable in relation to firms' financing capability.[1] Our contribution is, however, more general, as we explore optimal government intervention in a production economy with externalities, financing constraints and private information about the individual costs of reducing these externalities.

According to Pigou (1920), the optimal marginal tax on a good that generates externalities should be equal to the marginal social cost that arises from consumption or production of that good, evaluated at the respective (optimal) allocation. In this article, we take a linear tax cum transfer as the starting point and show that, in the presence of financing constraints, the optimal choice for the marginal tax is strictly below the respective Pigouvian level. This has the following two reasons. The first follows immediately from our introduction of financing constraints, which we derive from first principles in a model of moral hazard. The tax-induced necessity to raise a larger amount of financing, so as to thereby fund higher abatement costs, causes inefficiencies as it generates a greater wedge between the incentives of firm insiders

* Corresponding author: Roman Inderst, University of Frankfurt, and Imperial College London, Grüneburgplatz 1, 60323 Frankfurt am Main, Germany. Email: inderst@finance.uni-frankfurt.de.

We thank the editor Andrea Galeotti, three anonymous referees as well as seminar participants at Carlos III Madrid, Imperial College London and the University of Frankfurt for helpful comments.

[1] For instance, the UK government has recently set up a Green Investment Bank which will provide investment subsidies and low-interest loans. The UK Green Investment Bank plc was created in 2012 by the UK government for the financing of the private sector's investments related to environmental preservation and improvement. This was partly modelled around the German state-owned bank KfW, which among other things provides large-scale subsidised credit to businesses that apply energy-saving technologies or invest to reduce CO_2 output.

and the interests of (external) investors holding claims to the firm. The optimal linear tax takes these productive inefficiencies into account. A second reason for why the optimal linear tax is still lower arises when the concerned firms have different abatement costs. Then, as the tax on externalities increases, this affects more-than-proportionally firms with higher abatement costs. As a consequence, such a linear tax cum transfer leads to a reallocation of resources from firms with higher abatement costs to those with lower abatement costs. Again, this is not inconsequential for aggregate welfare as it results in a higher incremental inefficiency (due to the moral hazard problem *vis-à-vis* outside investors) for firms that have already a higher external financing need due to higher abatement costs.

The combination of taxes on externalities, outside financing and firm heterogeneity generates scope for richer policies than a linear tax. Such policies increase aggregate welfare by addressing in addition the described inefficiency that arises from the tax's differential impact across firms with different resulting financing needs. In this article, we discuss two such policies and relate them to observed practice.

We first consider the scope for a non-linear tax. In order to mitigate the tax-induced reallocation of resources, the optimal non-linear tax no longer prescribes a constant marginal tax on the externality. Instead, it is optimal to impose the highest marginal taxes for both relatively low and relatively high levels of pollution, thereby dampening the impact of the corrective tax for 'average polluters'. While this cannot be implemented by a scheme of tradable 'pollution rights' alone, a government could, to the extent that this is legally feasible, augment an existing supranational permit scheme by a subsidy that is paid per unit of required permit. In a coarse interpretation of our optimal non-linear tax, the subsidy should then be paid only over an interim range, i.e. for any required pollution permit above x 'units', but with a cap at y 'units'. Such a policy would alleviate the burden of costs from abatement for the 'average polluter', thus improving resource allocation and increasing aggregate welfare.

In particular, when the use of such non-linear taxes is restricted, for example, as it is not legally feasible to compensate the implications of a supranational system of tradable pollution rights in this way, resource allocation can be improved by introducing loan-based grants. Though these must also respect incentive compatibility, i.e. prevent opportunistic behaviour by firms that actually have lower abatement costs, they still dampen the effect of higher financing needs for firms with high abatement costs, thereby increasing aggregate welfare. Combining taxes on externalities with grants linked to loans, as is frequently observed in practice (discussed above), thus dominates taxes or pollution permits alone. It does so by providing the government with an additional instrument so as to thereby address in an incentive compatible way the differential impact that a (higher) tax on externalities has on firms with altogether lower or higher external financing requirements due to different abatement costs.

We consider as our main modelling contribution the introduction of financing constraints and thereby the need to raise external financing into a model of optimal taxation of externalities. More broadly, this relates our article to the literature that analyses the effect of liability on environmental care. In some of this literature (see the survey in Boyer *et al.*, 2006), compensation for damages is restricted by agents' limited resources or the limited liability embedded in the financial structure that they use to finance production. Imposing an extended liability also on the providers of outside finance, next to other contractual partners, may then impact efficiency, in particular in the presence of financial frictions and imperfect financial markets (Pitchford, 1995; Boyer and Laffont, 1997; Hiriart and Martimort, 2006; Tirole, 2010).[2]

As noted above, our results deviate from the Pigou rule, which would prescribe to set the tax so as to internalise the marginal social damage from pollution fully. Starting with Sandmo (1975), the public finance literature has also explored how the Pigou rule needs to be modified in a second-best setting where additional distortionary taxes are in place, mostly for government revenue raising purposes. One of the main results of this literature is the so-called 'additivity property' according to which optimal taxes can be expressed as the sum of the optimal Pigouvian tax and the optimal taxes in a problem without externalities as long as it is possible to tax directly the externality generating commodity.[3] Hence, the optimal pollution tax typically differs from the Pigouvian level.[4] We also look at optimal taxation of externalities in a second-best setting where firms are financially constrained. Still, in our model, the only reason for governmental intervention is to reduce externalities and the associated tax generates additional productive inefficiencies due to the moral hazard problem *vis-à-vis* outside investors, which are exacerbated in the aggregate when firms are heterogeneous. Notably, recent contributions in the public finance literature on mixed taxation problems with externalities have restored the Pigou rule, mostly through the use of optimal non-linear income taxes that compensate for tax-induced distortions on production and consumption in these models (cf. Cremer *et al.*, 1998; Jacobs and de Mooij, 2011; Kaplow, 2012).[5] We solve for the optimal mechanism and show that in the presence of outside financing constraints, there is still a wedge between the optimal marginal tax rate and the Pigouvian tax.

That divergences from Pigouvian corrections might be optimal in a second-best setting is also shown in Diamond (1973). In his model of consumption externalities, individuals differ in their sensitivity to the aggregate externality but the government is (exogenously) constrained to a uniform linear tax. When externalities are not separable, the second-best linear tax may deviate from the Pigouvian tax. This is the case as then a tax not only has a direct effect on the demand for the externality-causing good (leading to the Pigouvian component of the optimal correction) but also an indirect effect from the changes in consumption for the good induced by the direct effect (leading to an additional adjustment term).[6] With separable externalities, the

[2] The interaction between private financial frictions and public policy has been addressed also in the literature on entrepreneurship that examines various rationales for policy intervention, in particular the possible spillover effects created by start-ups. Boadway and Tremblay (2004) offer a broad overview of the literature, which mainly focuses on tax considerations.

[3] See Sandmo (1975) and more recently Kopczuk (2003).

[4] See for instance Bovenberg and de Mooij (1994) as well as, for a recent overview, Goulder and Bovenberg (2002) and the references therein.

[5] Cremer *et al.* (1998) and Jacobs and de Mooij (2011) show that no correction of the Pigouvian rule is needed with weakly separable preferences, taking into account distortions arising from redistributive income taxation. As pointed out in Kaplow (2012), the key is to recognise that the environmental tax will induce, *ceteris paribus*, redistribution but that the income tax schedule can be adjusted such as to neutralise this distributional effect completely. For earlier work in a similar vein see Diamond and Mirrlees (1971), who state that distributional concerns do not justify violating production efficiency if the government can optimally adjust taxes (on consumption).

[6] We thank an anonymous referee for pointing this out to us. In Diamond's model, these higher round equilibrium effects can thus lead both to an overcorrection and an undercorrection relative to the Pigouvian tax.

optimality of the Pigouvian tax is restored. Rothschild and Scheuer (2014*a*) consider a general framework of optimal non-linear income taxation in a multiactivity economy with externalities. Because, in their model, taxes affect the relative returns to different activities, they induce shifts in the allocation of effort across activities and the optimal tax can again be above or below the Pigouvian level.[7]

The rest of this article is organised as follows. Section 1 introduces the economy. Section 2 derives some preliminary results. In Sections 3 and 4, we solve for the optimal linear and non-linear tax on the externality. In Section 5, we allow the government to use, as an additional instrument, a grant linked to the size of firms' loans. Section 6 summarises our results. All proofs are contained in Appendix A. Some extensions as well as additional technical material to which we refer to in the main text are collected in the online Appendices.

1. The Economy

We consider an economy populated by a unit mass of agents indexed by $i \in I = [0, 1]$. There are two points of time: $t = 0$ and $t = 1$. Agents have access to the same production technology that pays out in the final period $t = 1$. Abstracting first from both the presence of a policy maker and the presence of externalities, starting production in $t = 0$ requires the investment of $I_0 \geq 0$ and generates in $t = 1$ to either zero output or an output of $x > 0$.[8] We employ the most standard model of moral hazard and thus stipulate that the likelihood of a positive outcome depends on the non-observable, real-valued effort $e \geq 0$ that the respective agent exerts. More precisely, effort e affects the probability $p(e)$ of high output at private disutility $c(e)$. Here, $c(e)$ and $p(e)$ are both (strictly) increasing and it is convenient to stipulate the following: $c'' > 0, p'' \leq 0, c'(0) = 0, p'(0) > 0$, and that $c'(e)/p'(e)$ becomes sufficiently large as $p(e)$ approaches one. Taken together, these conditions ensure a unique, interior solution for the effort choice problem of the agent.

Agents have originally zero resources and thus have to raise capital from investors located outside the considered economy in order to start production and cover all operating expenses. We stipulate that the agents' utility is linear in the resources that they consume and that they do not discount future consumption, which is why in our model, the only rationale for borrowing (i.e. raising outside finance) is for production. In terms of contracting with outside investors, we assume that the output realisation is verifiable and can thus be part of a financial contract.

1.1. *Externalities and Their Avoidance*

Production generates negative externalities, which can be reduced by additional investment. Precisely, we stipulate that when the activity of agent i creates $y_i \geq 0$ units of

[7] See also Rothschild and Scheuer (2014*b*) for a special case where agents can earn their income in traditional activities as well as through socially unproductive rent-seeking. When the government cannot observe the shares of individual agents' income earned in the two activities, the corrective component of the optimal non-linear income tax scheme in their model deviates from the Pigouvian correction due to a 'sectoral shift effect' in general equilibrium.

[8] In our model, the investment outlay and output are both measured in the same unit of 'resources' or capital.

these externalities, then this affects all other agents equally and, thereby, generates the social loss ρy_i with $\rho > 0$. We assume that utility is additively separable in these externalities. That is, when after $t = 1$ agent i is left with w_i resources for consumption and has exerted effort e_i, then his total utility is:

$$u_i = w_i - c(e_i) - \rho \int_{j \in I} y_j \mathrm{d}j. \tag{1}$$

Note that this implies that an agent's private incentives to reduce his own externality y_i are zero, given that the impact is distributed uniformly across all agents (with mass one). Externalities are verifiable.

It is now convenient, albeit this is without loss of generality, to stipulate that there is a given maximum level of externalities $\bar{y}$ (per firm). Consequently, the respective avoided externalities can be denoted by $a = \bar{y} - y$. Avoiding externalities, e.g. by using the respective technology mix or by operating production accordingly, is associated with a particular production cost.

An agent's cost of avoiding externalities depends on his type. Define the strictly positive real-valued type by θ_i, where we assume that θ_i is, for all $i \in I$, independently and identically distributed according to the distribution function $F(\theta)$, permitting a density function $f(\theta) > 0$ for all $\theta \in [\underline{\theta}, \bar{\theta}]$. We capture abatement costs by the twice continuously differentiable function $K(a, \theta)$ with $K(0, \theta) = 0$. Letting subscripts indicate partial derivatives with respect to the respective argument, we stipulate that $K_1(0, \theta) = 0$, $K_1(\bar{y}, \theta) = \infty$ and $K_{11}(a, \theta) > 0$, as well as $K_2(a, \theta) < 0$ for $a \in (0, \bar{y})$. Further, we assume that types θ are ordered such that:

$$K_{12}(a, \theta) := \frac{\partial^2 K(a, \theta)}{\partial a \partial \theta} < 0. \tag{2}$$

That is, higher types θ have everywhere strictly lower marginal costs of abatement. For instance, we could take $K(a, \theta) = k(a)/\theta$, where $k(a)$ inherits the properties of $K(a, \cdot)$. Note that the respective costs are incurred, together with the investment I_0, right when production starts in $t = 0$.

Our chosen set-up, where the reduction in negative externalities is a function of investment, allows also for the following alternative interpretation. We could think of y or likewise $a = \bar{y} - y$ as a (continuous) verifiable technology choice, for example, the 'amount' of additionally installed fuel-efficient equipment or of energy-efficient building material that is used when setting up production. Though the purchase costs may be the same for all agents, agents' costs of installation or, more generally, their total opportunity costs may differ, given the buildings, equipment and technologies that they already own. For higher θ, the associated costs are lower according to condition (2).[9]

[9] When taking this interpretation, a given choice of technology could then be associated with some (possibly stochastic) generation of externalities (for which we would then need a different notation). Without loss of generality, any policy could then, however, target directly the adoption of the technology, a. Also, the agents' utility function (1) and thus also the policy maker's objective function could be rewritten accordingly, namely as a function of expected externalities, without changing results.

1.2. *Feasible Policies*

We introduce a utilitarian policy maker, who maximises the expected utility of all agents: $\mathrm{E}[\int u_i \mathrm{d}i]$. For simplicity, we refer to the policy maker as the government and consider various policy instruments. Note that, if firms are not financially constrained or, likewise, there is no agency problem *vis-à-vis* outside investors, then the first-best can be implemented through a linear tax on externalities where the marginal tax is just equal to the marginal externality that arises from production (ρ) – the Pigou rule. Our benchmark, hence, is that of a linear tax on externalities, coupled with a transfer that is paid out of tax receipts. We characterise the optimal linear tax and show how with financial constraints this is strictly different from the Pigouvian tax. As we argue, the outcome can also be implemented through a market for pollution rights.

We then analyse various instruments that achieve a more efficient resource allocation than the optimal linear tax. Note first that there is no scope for the government to raise finance on behalf of agents, unless it would use this to, at the same time, reallocate resources. This can, however, also be achieved through providing grants linked to the amount of outside finance that agents privately raise, which is a policy that is frequently observed in practice (see the Introduction). We characterise the optimal grant scheme. With such grants in place, we argue further that there is no additional role that taxes levied on output could play for the purpose of efficiently reallocating resources.[10] A further improvement of efficiency can, however, be achieved when the linear tax – or, likewise, a market for pollution rights that induces such a linear tax – is replaced by a non-linear tax scheme. Here, we use a mechanism design approach to solve for the optimal such non-linear tax scheme and we illustrate the difference to the optimal linear tax with the help of numerical examples.

2. Preliminary Result: The Outside Financing Problem

Consider the problem of an agent with no initial funds. As we specify shortly, to start production, the agent thus needs external funding to cover an original investment outlay as well as abatement costs and to pay taxes on externalities. We denote the respective funding requirement, which is endogenised subsequently, by L. As the agent can only pay back in case output is positive, the contract with outside investors can be restricted to a single variable: The repayment R is made in case the output equals x. Given some repayment R, note that the agent's payoff is:

$$p(e)(x-R)-c(e),$$

so that the uniquely optimal effort level e^* is given by the first-order condition:

$$p'(e^*)(x-R)-c'(e^*)=0. \tag{3}$$

With the optimal effort choice determined from (3), denote the agent's surplus for given R by:

$$\omega=p(e^*)(x-R)-c(e^*).$$

[10] As we discuss below, the right to these tax receipts could then be sold *ex ante* so as to alleviate financial constraints (at least for some types).

Now, assuming either a competitive market for funds or that the agent makes a take-it-or-leave-it offer, for any given level of L, there is a unique repayment requirement R^* that maximises ω subject to (3) and the, by optimality, binding break-even requirement:[11]

$$p(e^*)R^* = L. \tag{4}$$

After substituting for this, we can write the corresponding maximum total surplus net of the respective funding expenditure as:

$$\omega(L) = p(e^*)x - c(e^*) - L. \tag{5}$$

LEMMA 1. *The surplus function $\omega(L)$ satisfies the following property due to the underlying agency problem between the respective firm and investors: For all feasible $L'' > L' \geq 0$, we have:*

$$\omega(L') - \omega(L'') > L'' - L'. \tag{6}$$

Proof. See Appendix A.

With $L = 0$ and thus $R^* = 0$, the agent would choose a first-best effort level, solving $p'(e_{FB})x - c'(e_{FB}) = 0$, thereby realising a total surplus of $\omega(0)$. Note that this is gross of externalities and all possible transfers. When, however, outside financing must be raised, so that investors must be promised a repayment $R^* > 0$ to ensure that they break even, the corresponding effort is from (3) strictly lower and thus only second best: $e^* < e_{FB}$. Moreover, the higher is the respective financing need and thus also the repayment that must be promised to (outside) investors, the larger is the wedge between the (inside) agent's incentives to provide effort and thus e^*, on the one hand, and the level of effort that would maximise the value of the whole firm, e_{FB}, on the other hand. The respective loss in surplus is captured by (6). At points where $\omega(L)$ is differentiable (which it is almost everywhere), this becomes $\omega'(L) < -1$: a (marginal) increase in external financing that is needed reduces total surplus by more than the respective outlay. Again, this captures the key (productive) inefficiency that arises from the agency problem due to non-observable effort and the need to raise outside finance, which will be one of the driving forces behind our results.

To streamline the subsequent exposition, we stipulate in what follows that $\omega(L)$ is (also twice) continuously differentiable everywhere. Consider, for instance, $p(e) = e$ and $c(e) = e^2/(2\gamma)$, where γ is taken to be sufficiently small so as to ensure that $p(e^*) = \gamma(x - R^*) < 1$ holds in equilibrium. This specification will also be used subsequently for some numerical illustrations. Then, $\omega(L)$ is smooth and, as is easily confirmed, indeed strictly decreasing for $L > 0$ with $\omega'(L) < -1$. Also, another

[11] In case (3) and (4) have a unique solution (e^*, R^*) for a given level of L, the optimal contract is determined from the two constraints alone. If, on the other hand, (3) and (4) allow for multiple solutions, the agent optimally picks the one maximising ω (see also the proof of Lemma 1).

important feature of the surplus function is easily established: $\omega''(L) < 0$.[12] That is, the larger becomes the outside financing need and thus the required repayment so as to make investors break-even, the larger is also the marginal negative impact on surplus that arises from the productive inefficiency due to the underlying moral hazard problem. In fact, the outside financing problem more generally has a (strictly) optimal solution in form of a deterministic contract for all L, if and only if $\omega(L)$ is strictly concave, which is what we assume in what follows.[13] A sufficient condition for this to be the case is that, at points of differentiability, it holds for the optimal effort that $\mathrm{d}^2e^*/(\mathrm{d}L)^2 < 0$.[14] When we take $p(e) = e$, as is common, this holds if the respective costs satisfy $c'''(e) \geq 0$.

3. Linear Tax and Tradable Pollution Rights

3.1. *Preliminary Analysis*

In this Section, we consider the problem of a utilitarian policy maker choosing the parameters of a linear tax policy. That is, depending on the volume of produced externalities, y_i, each agent is taxed according to the function:[15]

$$\tau(y) = \tau_0 + \tau_1 y. \tag{7}$$

Here, τ_1 is the per-unit tax on the externality, while the fixed component τ_0 takes into account the overall distribution that is achieved by (optimally) making the government's budget just balance:

$$\tau_0 + \tau_1 \int_{i \in I} y_i \mathrm{d}i = 0. \tag{8}$$

Our motivation for the restriction to such a linear tax is the following. First and foremost, as we argue in more detail below, such a scheme corresponds to the implementation of a system of tradable pollution rights. In that case, the government's choice parameter would be the aggregate volume of externalities. Taking as a benchmark the outcome where such a system of tradable pollution rights is in place, we later argue how this can be optimally complemented with additional policies, such as tax-subsidised loans. Further, the characterisation of the optimal linear tax will make

[12] With the restriction $x^2 - 4/\gamma L \geq 0$, which ensures financial feasibility for given L, the surplus function is given by:

$$\omega(L) = \frac{\gamma}{8}\left[x + \left(x^2 - \frac{4}{\gamma}L\right)^{\frac{1}{2}}\right]^2,$$

and it is easily verified that $\omega'(L) = -\frac{1}{2}\{1 + x[x^2 - (4/\gamma)L]^{-\frac{1}{2}}\} < -1$ and $\omega''(L) = -(x/\gamma)\left[x^2 - 4(L/\gamma)\right]^{-3/2} < 0$.

[13] Put differently, if $\omega(L)$ is (strictly) concave, then *ex ante* randomisation is (strictly) undesirable for all L, while, if $\omega(L)$ is (strictly) convex on some subset, then there exists a value L', such that a stochastic contract would be (strictly) desirable at L'. This result is well known in the contracting literature, see for instance Arnott and Stiglitz (1988).

[14] From (5) we immediately obtain $\omega''(L) = [p''(e^*)x - c''(e^*)](\mathrm{d}e^*)/(\mathrm{d}L)^2 + [p'(e^*)x - c'(e^*)](\mathrm{d}^2e^*)/(\mathrm{d}L^2)$, where the first term is clearly negative. The result then follows from $e^* < e_{FB}$, which implies that $p'(e^*)x - c'(e^*) > 0$.

[15] Strictly speaking, the tax schedule is an (affine) two-part tariff.

transparent how both the presence of financing constraints and agent heterogeneity generally affect the implications of taxing externalities and, thereby, the optimal level and form of government intervention.

Without loss of generality, we stipulate that the agent must 'purchase' the respective pollution rights (or pay the tax) when starting production in $t = 0$. Consequently, without resources on his own, an agent must raise outside finance equal to:

$$L(y, \theta) = \max\{I_0 + K(\bar{y} - y, \theta) + \tau(y), 0\}. \tag{9}$$

Note that, from $\omega'(L) < -1$, it is not optimal for the agent to raise outside finance for consumption while it is equally optimal to use all of his own resources to reduce the amount of external financing raised.

Given the tax scheme $\tau(y) = \tau_0 + \tau_1 y$, we consider first the programme of an individual agent. The agent chooses y_i and, consequently, has to raise $L(y_i, \theta_i)$, as given by (9). Dropping the subscript i, an agent of type θ thus chooses y to maximise $\omega[L(y, \theta)]$ with $L(y, \theta)$ given by (9).

Lemma 2. *Suppose the government imposes a (budget-balancing) tax-cum-transfer $\tau(y) = \tau_0 + \tau_1 y$. Then, an agent of type θ chooses the optimal level of externalities $y^*(\theta)$ and thus a unique level of abatement $a^*(\theta) = \bar{y} - y^*(\theta)$ so that*:

$$K_1[a^*(\theta), \theta] = \tau_1, \tag{10}$$

from which $a^*(\theta)$ is strictly increasing in both τ_1 and θ. Still, higher type agents incur lower abatement costs and thus need to raise less outside finance:

$$\frac{\mathrm{d}L(\cdot)}{\mathrm{d}\theta} = K_2[a^*(\theta), \theta] < 0. \tag{11}$$

Proof. See Appendix A.

Hence, with a linear tax on externalities, each agent chooses a level of abatement so that the marginal financial benefits that follow from a reduction in the incurred tax are equal to the marginal cost of abatement. Importantly, productive efficiency, as expressed through the slope $\omega'(\cdot)$, plays no role in this trade-off. Moreover, note that, under the agent's optimal choice, his need to raise outside finance is always strictly decreasing in his type. In fact, as the agent chooses his privately optimal level of abatement, depending on θ, this follows immediately from optimality, as otherwise higher type agents could not enjoy a higher expected utility $\omega(L)$.

3.2. *Optimal Tax*

The objective function of a utilitarian government is to maximise the expected utility of all agents:

$$\mathrm{E}(u_i) = \int_{i \in I} [\omega(L_i) - \rho y_i] \mathrm{d}i.$$

(This uses that $\omega(L_i)$ already takes into account the investment costs, as these are funded by outside investors.) Given the agent's optimal decision, using Lemma 2, the government's programme is then to choose τ_0 and τ_1 so as to maximise:[16]

$$\mathrm{E}(u_i) = \int_{\Theta} (\omega\{L[y^*(\theta), \theta]\} - \rho y^*(\theta))\mathrm{d}F(\theta), \tag{12}$$

subject to the budget-balancing constraint (cf. (8)):

$$\tau_0 + \tau_1 \int_{\Theta} y^*(\theta)\mathrm{d}F(\theta) = 0. \tag{13}$$

Take for a moment the benchmark without financial constraints, so that everywhere $\omega'(\cdot) = -1$. Then, from substitution of (13) into (12) while using the agent's first-order condition (cf. Lemma 2), we would obtain the Pigou rule $\tau_1 = \rho$. This obviously implements the first-best outcome, despite agents' private information about their marginal cost of abatement. The following result characterises, instead, the optimal linear tax when agents must raise outside finance and when this gives rise to a deadweight loss due to agency problems *vis-à-vis* outside investors.

PROPOSITION 1. *The optimal linear per-unit tax τ_1 satisfies*:

$$\begin{aligned}
&\tau_1 \times \left(-\int_{\Theta} \omega'\{L[y^*(\theta), \theta]\}\mathrm{d}F(\theta)\right) \\
&= \rho - \frac{\int_{\Theta} \omega'\{L[y^*(\theta), \theta]\}\left[y^*(\theta) - \int_{\Theta} y^*(\theta')\mathrm{d}F(\theta')\right]\mathrm{d}F(\theta)}{\int_{\Theta} \frac{\mathrm{d}y^*(\theta)}{\mathrm{d}\tau_1}\mathrm{d}F(\theta)},
\end{aligned} \tag{14}$$

which implies that τ_1 is strictly smaller than the marginal externality that arises from production (ρ).

Proof. See Appendix A.

From (14), an optimal tax τ_1 is strictly lower than ρ, the marginal social cost that arises from production. This is so for two (albeit related) reasons. The first reason is captured by the multiplier:

$$\left(-\int_{\Theta} \omega'\{L[y^*(\theta), \theta]\}\mathrm{d}F(\theta)\right) > 1 \tag{15}$$

on the left-hand side, using $\omega' < -1$; given $\int_{\Theta}(\mathrm{d}y^*(\theta)/\mathrm{d}\tau_1)\mathrm{d}F(\theta) < 0$, the second reason is captured by the term:

$$\left\{-\int_{\Theta} \omega'(\cdot)\left[y^*(\theta) - \int_{\Theta} y^*(\theta')\mathrm{d}F(\theta')\right]\mathrm{d}F(\theta)\right\} > 0, \tag{16}$$

which is subtracted on the right-hand side in (14), and which – as we argue – holds from $\omega'' < 0$. We now discuss both terms in turn.

[16] For given τ_1, $y^*(\theta)$ is determined from (10).

The term (15) captures the fact that, to reduce externalities, agents must raise outside finance. Due to the associated agency problem, this involves an additional 'shadow cost', namely in the form of lower productive efficiency as effort becomes inefficiently low. (Formally, $\omega'(\cdot) < -1$.) Next, the term (16) captures the efficiency implications of the reallocation of resources that goes hand-in-hand with the applied taxation, namely from agents with higher marginal abatement costs to agents with lower marginal abatement costs. The impact of such reallocation on aggregate productive efficiency is negative. This follows from the following two observations: first, with a linear tax, high-type agents incur, under the optimal choice $y^*(\theta)$, strictly lower costs of abatement; second, $\omega(\cdot)$ is strictly concave. As the tax on externalities shifts resources to high-type agents, this reduces the agency problem of high-type agents but increases the agency problem of low-type agents. Thus, it makes the already more productive high-type agents (endogenously) still more productive, while further reducing productivity of low-type agents. This leads to a reduction in aggregate efficiency of production in the economy.

3.2.1. *Pollution Permit Scheme*

It is straightforward to see that the government could implement the outcome of Proposition 1 also as follows. The government could set a total maximum capacity for externalities Y and allocate this uniformly (and for free) across all agents. Thus, each agent receives the same capacity, which we may write as $Y_i = Y$, as there is a unit measure of agents in the economy. These capacities (or 'pollution rights') are then traded in the market.

When τ_1 is the resulting price, we obviously have that $K_1[a^*(\theta), \theta] = \tau_1$, as previously in (10), together with:

$$\int_\Theta a^*(\theta)\mathrm{d}F(\theta) = \overline{y} - Y.$$

This uniquely links τ_1 to Y, and *vice versa*. The equivalence of the two policy instruments can then be seen immediately from substituting into the funding retirement (9):

$$\begin{aligned} L[y^*(\theta), \theta] &= I_0 + K[\overline{y} - y^*(\theta), \theta] + \tau[y^*(\theta) - Y] \\ &= I_0 + K[\overline{y} - y^*(\theta), \theta] + \tau\left[y^*(\theta) - \int_\Theta y^*(\theta)\mathrm{d}F(\theta)\right]. \end{aligned}$$

This is just the same as under the linear tax, after substituting the budget-balancing constraint for taxes (13).

Corollary 1. *The optimal linear tax can be implemented through a pollution permit scheme, where each agent receives the same level of tradable pollution rights Y.*

Observe as well that the allocation that is achieved under a linear tax is Pareto efficient. This observation follows immediately from the fact that the implemented abatement choices satisfy condition (10), so that the marginal costs of abatement are equalised across all agents. Under the respective permit scheme from Corollary 1, which leads to the same allocation, there is, thus, no scope for further Pareto improving 'trades'.

3.2.2. *Further Analysis*

While the allocation achieved under the optimal linear tax is Pareto efficient, in our model, financial frictions cause productive inefficiencies that can be reduced in the aggregate by employing more complicated policy instruments. In particular, Proposition 1 isolated two (related) reasons for why the optimal tax under financing constraints is strictly below the optimal Pigouvian tax: the shadow cost of raising financing, which is due to the agency problem, and the associated reallocation of resources, which in the aggregate exacerbates this agency problem. In the following, we discuss ways how policy makers could reduce these aggregate inefficiencies and, thereby, optimally induce a higher level of aggregate abatement at less (productive) inefficiency.

The derivation of the optimal non-linear tax in Section 5 further clarifies the tension between reducing externalities and reallocating resources, which in our case increases agency costs in the aggregate. We show that the optimal such non-linear tax would essentially 'dampen' the impact of the tax for 'average polluters', namely through imposing high marginal taxes for both relatively low and relatively high levels of pollution. As we show in a numerical example, however, the resulting efficiency gains may make it optimal to, thereby, inducing a strictly higher level of abatement for all agents.

Such a scheme of non-linear taxes may, however, not always be feasible, for instance, when a supranational pollution permit scheme restricts the scope of national policies. As we show, the government could then still increase social surplus through introducing grants that are linked to the amount of outside financing that each agent raises. As the agent's type is only privately known, these grants must still be incentive compatible to forestall opportunistic behaviour. With such grants in place, we argue further that there is no additional role that taxes levied on output could play for the purpose of efficiently reallocating resources. Still, for completeness, we show in online Appendix C that in situations where loan-based grants are not feasible, taxes on output can improve resource allocation and, thus, welfare.

4. Non-linear Taxes

4.1. *The Problem*

So far, we have restricted our analysis to a linear tax on externalities and identified the inefficiencies such a scheme causes when firms are financially constrained. First, the tax-induced additional abatement costs increase the required amount of external financing, thus exacerbating the agency problem *vis-à-vis* outside investors. Second, as firms have different (marginal) abatement costs, a tax-cum-transfer scheme leads to a reallocation of resources from high to low-cost types resulting in a further decline in average productive efficiency again working through the same agency problem. As we show below, it is this second inefficiency, caused by the reallocation of resources, that can be mitigated by relaxing the restriction to a linear tax on externalities. One way to counteract this resource reallocation, which is studied in this Section, is thus to allow the government to implement a general (non-linear) tax on externalities.

We thus depart from the assumption of a linear tax scheme and consider a general tax $\tau(y_i)$ as a function of the respective externalities y_i that agent i produces. Our approach is the following: instead of solving directly for the optimal non-linear tax, we set up a general mechanism-design problem. (This is then extended in the subsequent Sections to introduce other policy instruments.) Such a mechanism maps agents' truthful revelation of their type θ into both a prescribed level of externalities $y(\theta)$ and a respective transfer $T(\theta)$, which can be positive or negative. Once we have derived the optimal mechanism, we obtain from this the respective optimal tax scheme $\tau(y_i)$.

The mechanism must ensure truthtelling and thus incentive compatibility for each agent type.[17] That is, for all types $\theta \in \Theta$, it must hold that:

$$\omega\{L[y(\theta), T(\theta), \theta]\} \geq \omega\{L[y(\hat{\theta}), T(\hat{\theta}), \theta]\} \text{ for all } \hat{\theta} \in \Theta, \tag{17}$$

where:

$$\begin{aligned} L[y(\theta), T(\theta), \theta] &= I_0 + K[\bar{y} - y(\theta), \theta] + T(\theta), \\ L[y(\hat{\theta}), T(\hat{\theta}), \theta] &= I_0 + K[\bar{y} - y(\hat{\theta}), \theta] + T(\hat{\theta}). \end{aligned}$$

In other words, type θ does not strictly prefer to pretend to be any other type $\hat{\theta}$.

As is standard, we will employ optimal control techniques in solving the following optimisation problem and, therefore, restrict the mechanism $\{[y(\theta), T(\theta)]\}$ to piecewise continuously differentiable functions.[18] The incentive constraint (17) holds locally if 'truthtelling', i.e. $\hat{\theta} = \theta$, solves the respective first-order condition:

$$\left.\frac{\mathrm{d}\omega(L[y(\hat{\theta}), T(\hat{\theta}), \theta])}{\mathrm{d}\hat{\theta}}\right|_{\hat{\theta}=\theta} = \omega'(\cdot)\{T'(\theta) - y'(\theta)K_1[\bar{y} - y(\theta), \theta]\} = 0. \tag{18}$$

We presently assume that the 'first-order approach' is valid, so that (18) is sufficient to ensure global incentive compatibility. As is immediate from the single-crossing property (2), note that this requires also that the characterised function $y(\theta)$ be non-increasing.[19]

Define now (with some abuse of notation) the payoff function under truthtelling:[20]

$$U(\theta) = \omega\{L[y(\theta), T(\theta), \theta]\}.$$

We know that by incentive compatibility $U(\theta)$ is non-decreasing and continuous and thus almost everywhere continuously differentiable with:

$$\frac{\mathrm{d}U(\theta)}{\mathrm{d}\theta} = \left.\frac{\partial\omega(L(y[\hat{\theta}), T(\hat{\theta}), \theta])}{\partial\theta}\right|_{\hat{\theta}=\theta} = \omega'(\cdot)K_2[\bar{y} - y(\theta), \theta] > 0, \tag{19}$$

when $y(\theta) > 0$.[21]

[17] As is well known, under the considered environment, the restriction to direct, truthtelling mechanisms follows without loss of generality from the 'revelation principle'.

[18] Note that incentive compatibility alone requires that $y(\theta)$ has to be non-increasing and, hence, differentiable almost everywhere.

[19] In a previous version of the article, which is available from the authors upon request, we considered a more general solution ('second-order approach') allowing for the possibility of 'bunching'.

[20] This still presumes that taxes and subsidies are fully 'used' to increase or reduce the amount of funds that must be raised externally (instead of being immediately consumed or saved for consumption in the final period). As agents are not impatient and have risk neutral preferences, this restriction is without loss of generality.

[21] The expression in (19) follows from the envelope theorem, in particular, from the fact that the agent's optimal (truthful) report satisfies (18).

To solve for the optimal menu, we set up the government's optimal control problem. Define the financing requirement under truthtelling:

$$L(\theta) = I_0 + K[\bar{y} - y(\theta), \theta] + T(\theta),$$

which we take as the state variable. As thus $U(\theta) = \omega[L(\theta)]$, we have from (19) that:

$$\frac{\mathrm{d}L(\theta)}{\mathrm{d}\theta} = K_2[\bar{y} - y(\theta), \theta] < 0. \tag{20}$$

Further, from:

$$T(\theta) = L(\theta) - \{I_0 + K[\bar{y} - y(\theta), \theta]\}, \tag{21}$$

we can substitute pointwise for $T(\theta)$, once the state variable $L(\theta)$ as well as $y(\theta)$ have been determined. This leaves us with the single control variable $y(\theta)$.

Summing up, the government's objective is thus to maximise total utility over all agents:

$$\int_{\Theta} \{\omega[L(\theta)] - \rho y(\theta)\}\mathrm{d}F(\theta), \tag{22}$$

subject to the 'law of motion' (20) and the budget balance condition:

$$\int_{\Theta} \{L(\theta) - K[\bar{y} - y(\theta), \theta] - I_0\}\mathrm{d}F(\theta) = 0, \tag{23}$$

where we have substituted from (21).

4.2. *Characterisation*

We now relegate to the Appendix A the formulation of the respective Hamiltonian and the solution of the control problem. There, we also translate the solution into the optimal tax schedule. This is obtained from the characterised menu through substituting $\tau(y) = T[\theta(y)]$, where we use $\theta(y) = y^{-1}[y(\theta)]$.[22] Denote now the lowest and highest realised level of externalities by:

$$y_l = y(\overline{\theta}) < y_h = y(\underline{\theta}).$$

Proposition 2. *Under the optimal non-linear tax $\tau(y)$, the marginal tax rate is strictly positive, but strictly less than the marginal externality that arises from production (ρ), and highest at the two extremes, y_l and y_h. This implies that when the generated externality is already low (and abatement thus high), then the marginal benefits from further reducing pollution are strictly increasing. Instead, when the externality generated by the agent is still high (and abatement thus low), the marginal benefits from limiting pollution are strictly decreasing.*

Proof. See Appendix A.

While the inefficiency caused by the need to raise outside financing implies that $\tau'(y) < \rho$ for all levels of externalities, it holds from Proposition 2 that the optimal non-linear tax rewards a reduction of externalities in particular at very high and very low realisations, i.e. at the 'first units' and the 'last units'. The intuition for this result is as

[22] Note for this that we assume that the optimal mechanism prescribes a strictly decreasing level of externalities, with $y'(\theta) < 0$ (see footnote 19).

follows: at the heart is the attempt to restrict the reallocation of resources that is made to high-type agents as a consequence of the tax on externalities.[23] Recall that more resources are reallocated the higher the marginal tax. Obviously, at the lower boundary $\underline{\theta}$, there is no need to further distort the implemented choice of externalities, as there is 'no one' contributing to reallocation below $\underline{\theta}$. At the other extreme, when $\theta = \bar{\theta}$, there is also clearly no longer a benefit from further reducing the marginal tax since there is nobody benefiting from reallocation above $\bar{\theta}$.[24] These observations explain the derived properties of the tax scheme $\tau(y)$ at low and high levels of the externality.

One immediate implication of Proposition 2 is that a constant marginal tax alone is not optimal, at least not in our model. The reason is that it implies too much reallocation of resources away from firms with high abatement costs, which in our model increases aggregated deadweight loss in the economy. As discussed in the Introduction, if a system of tradable pollution rights is in place, for example, due to an international or supranational agreement, then Proposition 2 would suggest that an additional, non-linear scheme of taxes and subsidies, depending on the realised externality, would improve efficiency. To achieve the benefits of the optimal non-linear tax, it should essentially dampen the impact on the 'average polluter' but increase the impact on high-level and low-level polluters.

Our obtained characterisation for a continuum of types prescribes a marginal tax rate that varies constantly as we change the respective level of externalities y. A coarser 'practical' implementation of this may, instead, prescribe much less variation. For instance, along the discussed principles, a policy may then augment a mechanism of tradable pollution permits with a subsidy that is paid per unit of required permit but only over an interim range, for example, a subsidy of £3 for any required pollution right above 10,000 'units' but with a cap at, say, 20,000 'units'. That such a policy would alleviate the burden of costs from abatement for the 'average firms' would be an intended consequence, and the rationale for this would follow from our formal derivations.

4.3. *Illustration*

In order to gain more intuition about the characterisation of the optimal non-linear tax scheme, suppose now that $p(e) = e$, $c(e) = e^2/2\gamma$, and $K(a) = 1/\alpha\, a^2/\theta$, with θ distributed uniformly on $[\underline{\theta}, \bar{\theta}]$ $(\bar{\theta} > \underline{\theta} > 0)$. Recall from Section 3 that the agent's optimal effort in this case satisfies $e^* = \gamma(x - R)$. We suppose that γ is always sufficiently small so as to ensure an interior solution $e^* = p(e^*) < 1$ for the probability of success. Together with the investor's break-even condition, after solving for the equilibrium repayment requirement R, we obtain for the 'surplus function':

[23] This intuition is similar to findings in the optimal income taxation literature (Mirrlees, 1971; Seade, 1977, 1982), according to which the optimal marginal income tax is zero at the endpoints of the income scale (in the absence of bunching and with a bounded distribution of skills) and strictly positive elsewhere. There, the trade-off is between a redistributive gain and a negative incentive effect of a non-zero marginal tax.

[24] Recall that clearly also for the lowest and highest types the implemented choice of externalities will be below the first-best benchmark (or similarly the marginal tax rate strictly below ρ) due to the inefficiency caused by the need to raise external financing. However, as argued above, there is no need to further distort the level of externalities in order to reduce the reallocation of resources.

$$\omega(L) = \frac{\gamma}{8}\left[x + \left(x^2 - \frac{4}{\gamma}L\right)^{\frac{1}{2}}\right]^2,$$

which is, for $L > 0$, strictly decreasing with $\omega'(L) < -1$ and strictly concave. (The restriction that $x^2 - 4L/\gamma \geq 0$ ensures financial feasibility, for given L.) For this example, the respective solutions for the linear and non-linear taxes are derived numerically.

Figure 1 shows the marginal tax rate under the non-linear scheme, together with the optimal linear tax rate (left panel) as well as the difference in generated externalities under the optimal linear tax and the optimal non-linear tax scheme (right panel) for a particular specification: $\bar{\theta} = 8$, $\underline{\theta} = 1$, $x = 10$, $\alpha = 1$, $\gamma = 0.1$, $I_0 = 0.5$, and $\rho = 0.75$. We refer to this as a case with a (relatively) high degree of heterogeneity in types, as the difference $\bar{\theta} - \underline{\theta}$ is large compared to the case that we characterise further below. The optimal non-linear tax rate is U-shaped and maximised at the lowest and highest realised level of externalities, for which it is also strictly above the optimal linear rate. In this example, the optimal marginal tax rate at intermediate levels of externalities is, however, strictly below the optimal linear tax. As the second panel in Figure 1 shows, for intermediate types, the resulting level of externalities is also strictly lower under the optimal linear tax than under the optimal non-linear tax. This contrasts with our second example (Figure 2), where the marginal tax rate is strictly higher for all levels of externalities and where all types generate lower externalities compared to the case with the optimal linear tax. Compared to Figure 1, for Figure 2 we only change the low boundary of types from $\underline{\theta} = 1$ to $\underline{\theta} = 4$, thereby reducing the heterogeneity between agents and, thus, also the reallocation of resources generated by a (optimal) tax on externalities.

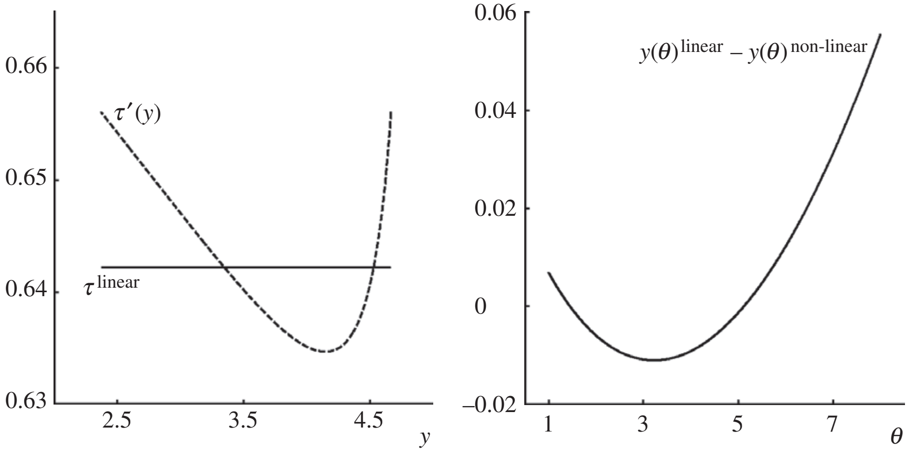

Fig. 1. *Example with High Degree of Heterogeneity* ($\bar{\theta} - \underline{\theta}$)
Notes. The left-hand side panel depicts the marginal tax rate under the optimal non-linear scheme and the optimal linear tax rate. The right-hand side panel shows the difference in generated externalities under the optimal linear and the optimal non-linear tax.

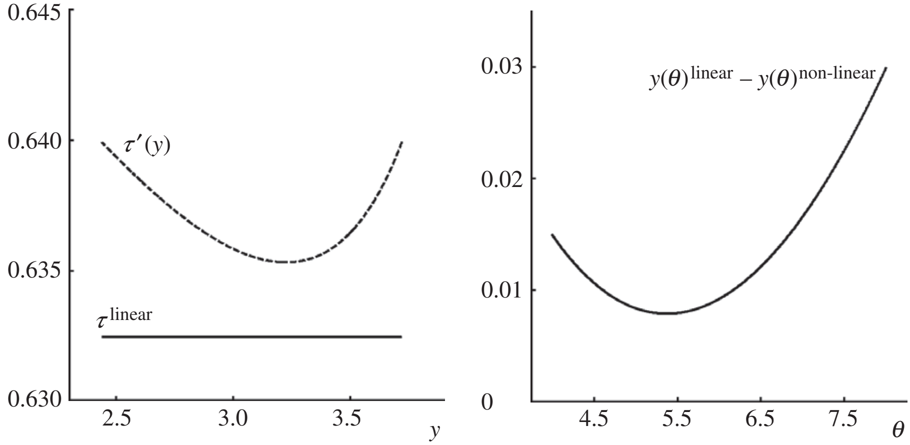

Fig. 2. *Example with Low Degree of Heterogeneity* $(\bar{\theta} - \underline{\theta})$

Notes. The left-hand side panel depicts the marginal tax rate under the optimal non-linear scheme and the optimal linear tax rate. The right-hand side panel shows the difference in generated externalities under the optimal linear and the optimal non-linear tax.

5. Loan-based Grants

So far, policy intervention has been restricted to a tax on externalities. As discussed in the Introduction, governments frequently use also grants linked to loans as a way to steer firm behaviour, in particular in the context of environmental policies.

We consider the following extension to our present model. We now allow government policy to be made contingent on the amount of financing that is raised by each agent. Importantly, note that the government cannot directly verify the real cost of abatement $K(\cdot)$. As noted above, even when a loan-based grant is tied to specific expenditures, e.g. for particular equipment, the overall costs, including opportunity costs, could still substantially differ between firms. In essence, what we use in what follows is the restriction that while the actual loan that a firm raises is verifiable, its true financial needs are still the firm's private information. Nevertheless, loan-based grants will prove effective in improving efficiency.

Suppose that agents are taxed on their externality according to a constant marginal rate τ_1. As noted above, our present restriction to such a linear tax is warranted in circumstances where non-linear taxes are not feasible due to the existence of a supranational system of tradable pollution rights.[25] With a linear tax, the optimal level of abatement for each agent satisfies the first-order condition:

$$K_1[\bar{y} - y^*(\theta), \theta] = \tau_1. \tag{24}$$

That is, irrespective of the firm's overall financial needs, including other taxes or grants, the firm simply chooses $y^*(\theta)$ to minimise expenditures. Taking for now τ_1 as given, an additional instrument thus serves the purpose of reducing the reallocation of resources that is generated by $\tau_1 > 0$. A loan-based grant (or tax) stipulates a positive

[25] Still, in online Appendix B, we also solve for the optimal mechanism when both non-linear taxes and loan-based grants can be used.

(or negative) payment $G(L_i)$, given an agent's loan volume L_i. As previously, it proves, however, convenient to first set up the problem in the language of mechanism design.

5.1. *Truthful Mechanism*

Such a mechanism now specifies for each agent a loan level $L(\theta)$ together with a payment made by each agent $t(\theta)$. It is immediate that under the optimal mechanism, $L(\theta)$ will just be sufficient to cover the agent's true expenditures. However, when an agent deviates and mimics another type $\hat{\theta}$ by raising a higher-than-necessary loan, then his pay-off becomes:[26]

$$U(\theta,\hat{\theta}) = \omega[L(\hat{\theta})] + \{L(\hat{\theta}) - I_0 - K[\bar{y} - y^*(\theta),\theta] - \tau_1 y^*(\theta) - t(\hat{\theta})\}. \tag{25}$$

Here, the term in square brackets captures the amount of financing that is raised above the true financial needs, which are unobservable to the government.[27] The mechanism is locally incentive compatible if $U(\theta,\,\hat{\theta})$ is maximised at $\hat{\theta} = \theta$, so that the additional term in (25) is zero: $U(\theta) = \omega[L(\theta)]$.

When $t(\theta)$ cannot condition on θ, as was the case in Section 3, then we have that:

$$\frac{\mathrm{d}U(\theta)}{\mathrm{d}\theta} = \omega'[L(\theta)]L'(\theta) = \omega'[L(\theta)]K_2[\bar{y} - y^*(\theta),\theta]. \tag{26}$$

This describes the slope of agents' utility and thus the extent to which a pure (linear) tax on externalities leads to a reallocation of resources in the economy.[28] (Clearly, when $y^*(\theta) = \bar{y}$ holds for all agents, then $U(\theta) = \omega(I_0)$ is constant.) We show now how by linking $t(\theta)$ to $L(\theta)$, this slope can be reduced, so that resources are more evenly distributed across agents.

We know from incentive compatibility that $U(\theta)$ must always be non-decreasing (in fact, it is strictly increasing when $y(\theta) > 0$). More precisely, this holds for any choice of $t(\theta)$ that is still incentive compatible. Therefore, the government wants to make the derivative $t'(\theta) > 0$ for all θ as steep as possible. (Recall that we frame $t(\theta)$ as a payment from agents, an analogous to our previous approach.) The upper boundary on $t'(\theta) > 0$ is obtained from the first-order condition of the agent's reporting problem.[29] Making use of (24), this yields for the slope of agents' pay-off:

$$\frac{\mathrm{d}U(\theta)}{\mathrm{d}\theta} = \omega'[L(\theta)]L'(\theta) = t'(\theta) - L'(\theta) = -K_2[\bar{y} - y^*(\theta),\theta], \tag{27}$$

which, for a given schedule $y^*(\theta)$ is clearly smaller than the slope in (26), showing the benefits of linking transfers to the amount of financing raised. Similarly, one obtains for the slope of the transfer $t(\theta)$:[30]

[26] Note that we have dropped any fixed part τ_0 for the tax on externalities, as this can be subsumed into $t(\theta)$.

[27] Note again, that we assume here that, for a firm of type θ, $y^*(\theta)$ is fixed from (24) independent of the report $\hat{\theta}$. Still, the results of this Section continue to hold if we, instead, assumed that a firm of type θ, reporting $\hat{\theta}$, would have to implement $y^*(\hat{\theta})$.

[28] The respective slope of agents' utility under non-linear taxation is given in (19), which differs from (26) due to the difference in the level of externalities generated under the two taxation regimes.

[29] More precisely, this only considers a feasible deviation to a type $\hat{\theta} < \theta$.

[30] An alternative way of expressing this is to substitute for $L'(\theta) = K_2[\bar{y} - y^*(\theta),\theta] + t'(\theta)$, so that $t'(\theta) = -K_2[\bar{y} - y^*(\theta),\,\theta]\{1 + \omega'[L(\theta)]/\omega'[L(\theta)]\}$.

$$t'(\theta) = L'(\theta)\{1 + \omega'[L(\theta)]\}. \tag{28}$$

Clearly, when $\omega'(\cdot) = -1$, as in the case without an agency problem of external financing, then $t'(\theta) = 0$. Otherwise, we have that $t'(\theta) > 0$. In fact, the strictly larger is $\omega'[L(\theta)] < -$ in absolute terms, the steeper can the loan-based grant be made, for the moment as a function of θ. (We translate this back into a function $G(L)$ below.) Such a transfer back to low-type agents, now linked to a loan, is made possible precisely as raising outside finance generates a deadweight loss. This holds also when more outside finance is raised than actually needed or, in a more general context, when a firm takes out a subsidised loan to fund abatement activities even though this is necessary, given the firm's resources. The deadweight loss of raising (too much) external finance prevents agents with lower costs of abatement from claiming a higher grant that is intended for agents with higher marginal costs of abatement.

5.2. *Optimal Loan-based Grant Scheme*

Under a grant scheme $G(L)$, we must have that:

$$G[L(\theta)] = -t(\theta). \tag{29}$$

Differentiating (29) and substituting from (28), we have that:

$$G'(L) = -\{1 + \omega'[L(\theta)]\} > 0, \tag{30}$$

which tells us how under the incentive compatible grant scheme, the absolute subsidy $G(L)$ varies with the loan size L.

PROPOSITION 3. *Suppose that a constant marginal tax τ_1 on externalities is in place. When the loan size L_i is verifiable, then the government can strictly increase efficiency by introducing, in addition, a loan-based grant. The optimal grant G(L) is strictly increasing in loan size with $G'(L)$ given by (30).*

Proof. See Appendix A.

Efficiency can thus be improved by linking transfers to the outside financing that agents raise to cover their abatement costs. By combining a tax on externalities with grants that are linked to loans, the reallocation of resources that is generated by the tax on externalities can be mitigated. Still, it is not profitable for agents with lower abatement costs to mimic those with higher abatement costs, so as to claim additional grants. This follows as raising more-than-necessary outside finance is costly as it exacerbates the agency problem *vis-à-vis* outside investors. As a consequence, combining taxes on externalities with grants linked to loans – as is frequently observed in practice – dominates, in our setting, pure linear taxes or pollution permits.[31]

In practice, even though economic instruments such as tradable permits or emission taxes are used to regulate emissions, they are often applied in

[31] As we show in online Appendix B, combining taxes with grants linked to loans continues to dominate a pure tax on externalities also when allowing for general non-linear taxes.

combination with other interventions. Frequently used additional instruments of support are subsidies to environmentally friendly investment. As noted in the Introduction, for instance, in Germany, a state-owned bank (KfW) provides on a large-scale subsidised credit to businesses that apply energy-saving technologies or invest to reduce CO_2 output. The UK government, in turn, is just about to set up a 'Green Investment Bank' which will provide investment subsidies and low-interest loans to accelerate private-sector investment in environmentally friendly infrastructure.

5.3. *Discussion*: *Taxes on Output*

As noted above, we have throughout restricted potential reallocation of resources to the initial stage. As we argue now, however, this is without loss of generality once we allow for loan-based grants. Consider thus taxes on output. The right to these taxes could then be sold *ex ante* so as to alleviate (at least for some types) the external financing constraint.[32] At first, it may seem that this gives the government an additional lever for reallocation of resources, as agents with lower abatement costs and thus lower financing needs end up with a higher equilibrium probability of success. However, the link from the agent's type to the likelihood of realising high *versus* low output is only indirect, namely through the agency problem and, therefore, through the amount of outside financing that is raised by the agent. A loan-based grant thus provides such a reallocation more directly through linking transfers to financing.[33] We show in online Appendix C how a tax on output could, however, increase efficiency when we assume that loan-based grants are not feasible.

6. Concluding Remarks

This article analyses the optimal policy towards externalities in the light of two constraints. First, agents who generate such (negative) externalities must raise outside finance when they want to increase their abatement. This generates inefficiencies due to agency costs *vis-à-vis* outside investors. Second, marginal abatement costs are private information, so that policies must be incentive compatible. We generate three sets of results in a simple, highly stylised model with these features. Our first result is that the optimal linear tax is strictly smaller than the benchmark Pigou tax, which would be equal to the marginal benefits from lowering externalities. In fact, we isolate two reasons for why this is the case: first, raising the necessary finance generates productive inefficiencies; second, with heterogeneous agents a higher tax generates aggregate productive inefficiencies as it leads to a reallocation of resources, thereby exacerbating aggregate financial frictions.

[32] We discuss these options explicitly in online Appendix C, where we consider the case of taxes on output without a loan-based grant.

[33] In this sense, in the presence of loan-based grants, output does not provide an additional 'tag' that would be optimally used for transfers. Originally, the term 'tagging' was coined by Akerlof (1978) to describe the use of taxes contingent on personal characteristics in order to improve on a purely income-based tax scheme; see. also Mankiw and Weinzierl (2010) for a recent application.

Our second result is that a non-linear tax on externalities enhances efficiency, as it allows the achievement of a given aggregate level of abatement more efficiently. As we show, this is the case as the non-linear tax allows to limit reallocation of resources to higher type agents with lower marginal abatement costs. We further show that under the optimal non-linear tax, the marginal benefits of abatement are highest for low and high levels of abatement (the 'first units' and the 'last units'). Importantly, this cannot be implemented by a scheme of tradable 'pollution rights'.

Our third result is that the government can further improve efficiency by linking transfers to the outside finance that agents raise so as to (purportedly) reduce externalities. In contrast to non-linear taxes, it can also be used when there is a (supranational) system of tradable pollution rights in place, which essentially implements a linear tax on externalities. Though agents with lower abatement costs can still mimic those with higher abatement costs, when additional grants are linked to credit, this becomes more costly, simply as raising more-than-necessary outside finance exacerbates the agency problem *vis-à-vis* outside investors. As a consequence, using jointly taxes on externalities and grants linked to loans – as is frequently observed in practice – may be a useful instrument, as it allows to improve aggregate productive efficiency.

Appendix A. Omitted Proofs

Proof of Lemma 1. The proof is complicated by the fact that, for given L, (3) and (4) together may have multiple solutions, which is why the following proof is not based on implicit differentiation alone. Solely for the purpose of this proof, we introduce some additional notation. Denote more explicitly $\omega(R, L)$, where we have already substituted for the unique optimal effort level, $e^*(R)$. Implicitly differentiating (3) reveals that $\mathrm{d}e^*/\mathrm{d}R < 0$ holds strictly. Note that, holding L constant, $\omega(R, L)$ is clearly strictly decreasing in R, so that when (3) and (4) have multiple solutions, there is still a unique value $R^*(L)$ that maximises $\omega(R, L)$. (In fact, it is the lowest such value.) Denote thus $\omega(L) = \omega[R^*(L), L]$. We next argue that $R^*(L)$ is strictly increasing. We argue to a contradiction and suppose instead the following: $L'' > L'$ with $R'' = R^*(L'')$, $R' = R^*(L')$, and $R'' \le R'$. This is, however, not optimal, as then R' would fail to maximise $\omega(R, L')$. To see this, take L' as given and consider now a repayment of R''. Then from $L' < L'' = e^*(R'')R''$ the break even constraint (4) is slack such that R'' is feasible. But $R'' \le R'$ implies from (3) that $e^*(R'') \ge e^*(R')$, such that, taken together, $\omega(R'', L') > \omega(R', L')$, which contradicts the optimality of R' at L'.

Given that $R^*(L)$ is strictly increasing, we thus have also that the resulting effort level, which is compactly expressed as $e^* = e^*[R^*(L)]$, is strictly decreasing in L. Consider now $\omega(L)$ as given in (5) and note that $p(e^*)x - c(e^*)$ is uniquely maximised by some value $e_{FB} = e^*(0)$ and for all $e^* < e_{FB}$ strictly increasing in e^*. This establishes property (6). Note finally that there is a maximum feasible loan size L so that (3) and (4) can indeed be jointly satisfied.

Proof of Lemma 2. We have:

$$\frac{\mathrm{d}\omega}{\mathrm{d}y} = [\tau_1 - K_1(\bar{y} - y, \theta)]\omega'[L(y, \theta)].$$

As the problem is strictly quasiconcave, this yields the optimality condition (10). From implicit differentiation and using (2) we have further that:

$$\frac{\mathrm{d}a^*(\theta)}{\mathrm{d}\theta} = -\frac{K_{12}[a^*(\theta),\theta]}{K_{11}[a^*(\theta),\theta]} > 0,$$
$$\frac{\mathrm{d}a^*(\theta)}{\tau_1} = \frac{1}{K_{11}[a^*(\theta),\theta]} > 0.$$

Finally, expression (11) follows from substituting the first-order condition (10) into the total derivative of $L(\cdot)$.

Proof of Proposition 1. We can substitute from (13) to obtain for each type the financing requirement:

$$L[y^*(\theta),\theta] = I_0 + K[\bar{y} - y^*(\theta),\theta] + \tau_1 y^*(\theta) - \tau_1 \int_\Theta y^*(\theta')\mathrm{d}F(\theta'),$$

so that $\mathrm{d}E[u_i]/\mathrm{d}\tau_1$ equals:

$$\int_\Theta \left[\begin{array}{l} \left(\left[y^*(\theta) - \int_\Theta y^*(\theta')\mathrm{d}F(\theta')\right] + \left\{ \frac{\mathrm{d}y^*(\theta)}{\mathrm{d}\tau_1}[\tau_1 - K_1(\cdot)] - \tau_1 \int_\Theta \frac{\mathrm{d}y^*(\theta')}{\mathrm{d}\tau_1}\mathrm{d}F(\theta') \right\} \right) \times \omega'(\cdot) \\ -\rho \frac{\mathrm{d}y^*(\theta)}{\mathrm{d}\tau_1} \end{array} \right] \mathrm{d}F(\theta).$$

Substituting the first-order condition (10) for $y^*(\theta)$, $\tau_1 = K_1$, this gives rise to the first-order condition (14). From:

$$\int_\Theta \frac{\mathrm{d}y^*(\theta)}{\mathrm{d}\tau_1}\mathrm{d}F(\theta) < 0,$$

given that $y^*(\theta)$ is strictly decreasing, it remains to prove that:

$$\int_\Theta \omega'(\cdot)\left[y^*(\theta) - \int_\Theta y^*(\theta')\mathrm{d}F(\theta')\right]\mathrm{d}F(\theta) < 0. \tag{A.1}$$

To see this, note first that, next to $\mathrm{d}y^*(\theta)/\mathrm{d}\tau_1 < 0$, we have from (11) and the concavity of $\omega(\cdot)$ that:

$$\frac{\mathrm{d}}{\mathrm{d}\theta}\{\omega'[L(y,\theta)]\} = \frac{\mathrm{d}L(\cdot)}{\mathrm{d}\theta}\omega'' > 0. \tag{A.2}$$

Define now the unique type $\hat{\theta}$ satisfying:

$$y^*(\hat{\theta}) = \mathrm{E}(y^*) = \int_\Theta y^*(\theta')\mathrm{d}F(\theta'),$$

while $y^*(\theta) > \mathrm{E}(y^*)$ holds for $\theta < \hat{\theta}$ and $y^*(\theta) < \mathrm{E}[y^*(\theta)]$ holds for $\theta > \hat{\theta}$. We can now rewrite the left-hand side of (A.1) as:

$$LS = \int_{\theta < \hat{\theta}} \omega'(\cdot)[y^*(\theta) - \mathrm{E}(y^*)]\mathrm{d}F(\theta) + \int_{\theta > \hat{\theta}} \omega'(\cdot)[y^*(\theta) - \mathrm{E}(y^*)]\mathrm{d}F(\theta). \tag{A.3}$$

There, the terms in the first integral are all strictly negative and the terms in the second integral are all strictly positive. Given strict monotonicity of $\omega'(\cdot)$, we can thus derive the upper bound:

$$\begin{aligned} LS &< \int_{\theta < \hat{\theta}} \omega'[L(y,\hat{\theta})][y^*(\theta) - \mathrm{E}(y^*)]\mathrm{d}F(\theta) + \int_{\theta < \hat{\theta}} \omega'[L(y,\hat{\theta})][y^*(\theta) - \mathrm{E}(y^*)]\mathrm{d}F(\theta) \\ &= \omega'[L(y,\hat{\theta})]\left[\int_\Theta \left\{ y^*(\theta) - \int_\Theta y^*(\theta')\mathrm{d}F(\theta')\right]\mathrm{d}F(\theta) \right\} = 0. \end{aligned}$$

This implies that the right-hand side of (14) is strictly smaller than ρ, so that together with the preceding argument we have indeed that $\tau_1 < \rho$.

Proof of Proposition 2. The Hamiltonian is given by:

$$H = \{\omega[L(\theta)] - \rho y(\theta)\}f(\theta) + \eta\{L(\theta) - K[\bar{y} - y(\theta), \theta] - I_0\}f(\theta) + \lambda(\theta)K_2[\bar{y} - y(\theta), \theta].$$

An optimal solution must satisfy the first-order condition for $y(\theta)$:

$$f(\theta)\{-\rho + \eta K_1[\bar{y} - y(\theta), \theta]\} - \lambda(\theta)K_{12}[\bar{y} - y(\theta), \theta] = 0 \tag{A.4}$$

and for the costate variable:

$$\frac{\partial H}{\partial L} = -\lambda'(\theta) \Leftrightarrow f(\theta)\{\omega'[L(\theta)] + \eta\} + \lambda'(\theta) = 0. \tag{A.5}$$

There are no terminal conditions, and the transversality conditions are given by:

$$\lambda(\overline{\theta}) = 0, \tag{A.6}$$

$$\lambda(\underline{\theta}) = 0. \tag{A.7}$$

Using (A.6) and (A.7), we thus obtain from integrating (A.5):

$$\lambda(\theta) = \int_{\theta}^{\overline{\theta}} \{\omega'[L(\vartheta)] + \eta\}\mathrm{d}F(\vartheta) = -\int_{\underline{\theta}}^{\theta} \{\omega'[L(\vartheta)] + \eta\}\mathrm{d}F(\vartheta)$$

and

$$\eta = -\int_{\underline{\theta}}^{\overline{\theta}} \omega'[L(\vartheta)]\mathrm{d}F(\vartheta) > 1. \tag{A.8}$$

Here, η expresses the marginal benefits when the economy's resource constraint was marginally relaxed (e.g. by some initial endowment that could be allocated by the government).

From (20) and the concavity of $\omega[L(\theta)]$, we have that $\omega'[L(\theta)]$ is increasing in θ. Thus, making use of (A.8):

$$\lambda'(\theta) = f(\theta)\int_{\underline{\theta}}^{\overline{\theta}} \{\omega'[L(\vartheta)] - \omega'[L(\theta)]\}\,\mathrm{d}F(\vartheta),$$

is first positive and then negative, i.e. $\lambda(\theta)$ is first increasing and then decreasing. Clearly, the transversality conditions (A.6) and (A.7) then imply that $\lambda(\theta) \geq 0$ holds everywhere, which we use in what follows.

Rearranging now the first-order condition for the control $y(\theta)$ in (A.4), we have@

$$\eta K_1[\bar{y} - y(\theta), \theta] = \rho + \frac{\lambda(\theta)}{f(\theta)}K_{12}[\bar{y} - y(\theta), \theta], \tag{A.9}$$

which, using $K_{12}(\cdot) < 0$, $\lambda(\cdot) \geq 0$ and $\eta > 1$, implies first that:

$$K_1[\bar{y} - y(\theta), \theta] < \rho.$$

That is, also with non-linear taxes, externalities are for all types higher than under the Pigou tax. Moreover, note that it holds only at the boundaries $\underline{\theta}$ and $\overline{\theta}$ (when they are finite) that $\lambda(\theta) = 0$ and thus:

$$\eta K_1[\bar{y} - y(\underline{\theta}), \underline{\theta}] = \rho \text{ and } \eta K_1[\bar{y} - y(\overline{\theta}), \overline{\theta}] = \rho. \tag{A.10}$$

Instead, for all other types we have the following:

$$\eta K_1[\bar{y} - y(\theta), \theta] < \rho \text{ for all } \theta \in (\underline{\theta}, \overline{\theta}). \tag{A.11}$$

Hence, under the optimal mechanism the marginal abatement costs are highest at the lowest and at the highest type, when evaluated at the respective choice $y(\theta)$. We finally analyse how this translates into the respective non-linear tax scheme $\tau(y_i)$. For this the following observation is useful. We obtain:

$$\begin{aligned} T'(\theta) &= L'(\theta) - K_2(\cdot) + y'(\theta)K_1(\cdot) \\ &= y'(\theta)K_1(\cdot) \le 0. \end{aligned}$$

That is, the non-linear tax on the externality still involves a transfer from low-type agents to high-type agents, given that $T'(\theta) \le 0$ (and strictly so where $y'(\theta) < 0$).[34]

We now substitute $\tau(y) = T[\theta(y)]$, where we use $\theta(y) = y^{-1}[y(\theta)]$. Note for this that we assume that the optimal mechanism prescribes a strictly decreasing level of externalities, with $y'(\theta) < 0$. That is, while $y(\theta)$ must be non-increasing from incentive compatibility, there is also no 'bunching' (footnote 19). From substituting the obtained characterisation, we then have that:

$$\begin{aligned} \tau'(y) &= T'(\theta)\frac{d\theta}{dy} = \frac{T'(\theta)}{y'(\theta)} = K_1[\bar{y} - y, \theta(y)] \\ &= \frac{1}{\eta}\left\{\rho + \frac{\lambda[\theta(y)]}{f[\theta(y)]}K_{12}[\bar{y} - y, \theta(y)]\right\}, \end{aligned} \tag{A.12}$$

such that the marginal tax is always (strictly) positive but also strictly smaller than the Pigouvian tax.

Here, as discussed above, the term $\eta > 1$ (cf. expression (A.8)) applies to all types and creates a first wedge between the 'Pigou tax' and the marginal tax with outside financing and agency costs. Turn now to the second term in rectangular brackets. As $K_{12} < 0$ (the key 'sorting condition' (2)) and as we obtained that $\lambda(\theta) \ge 0$, this term is negative and now captures the second rationale for why the optimal marginal tax is strictly lower, namely the inefficient reallocation of resources that goes along with the tax on externalities in our model. Note, however, that $\lambda(\underline{\theta}) = \lambda(\bar{\theta}) = 0$ holds at the boundaries, where under the optimal non-linear tax this effect no longer plays a role (A.10). However, for all $\theta \in (\underline{\theta}, \bar{\theta})$ we have that $\lambda(\theta) > 0$.

From further differentiating, we obtain next:

$$\eta\tau''(y) = -\frac{\lambda(\theta)}{f(\theta)}K_{112} + \frac{d\theta}{dy}\left\{\frac{\lambda(\theta)}{f(\theta)}K_{122} + K_{12}\frac{d}{d\theta}\left[\frac{\lambda(\theta)}{f(\theta)}\right]\right\}. \tag{A.13}$$

Expression (A.13) describes how the marginal tax changes. Recall once more from the transversality conditions (A.6)–(A.7) that at the boundaries $\underline{\theta}$ and $\bar{\theta}$ (when they are finite) we have $\lambda(\theta) = 0$. Recall that in the main text, we have defined the lowest and highest realised level of externalities by $y_l = y(\bar{\theta}) < y_h = y(\underline{\theta})$. Further, recall that $\lambda' > 0$ for low and $\lambda' < 0$ for high θ, while $d\theta/dy < 0$ (when there is no 'bunching'). Using then:

$$\frac{d}{d\theta}\left[\frac{\lambda(\theta)}{f(\theta)}\right] = \frac{1}{f^2(\theta)}[f(\theta)\lambda'(\theta) - f'(\theta)\lambda(\theta)],$$

we have at the 'endpoints':

$$\begin{aligned} \eta\tau''(y_l) &= \frac{d\theta}{dy}K_{12}\frac{\lambda'(\bar{\theta})}{f(\bar{\theta})} < 0, \\ \eta\tau''(y_h) &= \frac{d\theta}{dy}K_{12}\frac{\lambda'(\underline{\theta})}{f(\underline{\theta})} > 0. \end{aligned}$$

[34] In fact, incentive compatibility implies that in both cases, i.e. with linear and non-linear taxes, the marginal tax w.r.t. the agent's type is given by $y'(\theta)K_1(\cdot)$. (For the linear tax we can use that $T'(\theta) = \tau_1 y'(\theta)$ and that $K_1(\cdot) = \tau_1$.)

In other words, under the optimal non-linear tax $\tau(y)$, at very high levels of y the marginal tax $\tau'(y) > 0$ becomes strictly decreasing, while at very low levels of y the marginal tax is strictly increasing. This completes the characterisation of the solution and the proof.

Proof of Proposition 3. What remains to be shown is that the characterised mechanism is globally incentive compatible. Hence, we need to show that, for all θ and $\hat{\theta}$ with $\theta > \hat{\theta}$, it holds that:

$$U(\theta) \geq U(\theta, \hat{\theta}).$$

Using (25), we can rewrite this inequality to obtain:

$$\omega[L(\theta)] - \omega[L(\hat{\theta})] \geq K[\overline{y} - y^*(\hat{\theta}), \hat{\theta}] - K[\overline{y} - y^*(\theta), \theta] + \tau_1 y^*(\hat{\theta}) - \tau_1 y^*(\theta).$$

From $U(\theta) = \omega[L(\theta)]$ together with (27), we then can write equivalently:

$$\begin{aligned}\int_{\hat{\theta}}^{\theta} \frac{\mathrm{d}U(\varphi)}{\mathrm{d}\varphi}\mathrm{d}\varphi &= \int_{\hat{\theta}}^{\theta} -K_2[\overline{y} - y^*(\varphi), \varphi]\mathrm{d}\varphi \\ &\geq -\int_{\hat{\theta}}^{\theta} \frac{\mathrm{d}K[\overline{y} - y^*(\varphi), \varphi]}{\mathrm{d}\varphi}\mathrm{d}\varphi - \tau_1 \int_{\hat{\theta}}^{\theta} \frac{\mathrm{d}y^*(\varphi)}{\mathrm{d}\varphi}\mathrm{d}\varphi,\end{aligned}$$

which can be simplified, using (24), to obtain:

$$\begin{aligned}&-\int_{\hat{\theta}}^{\theta} K_2[\overline{y} - y^*(\varphi), \varphi]\mathrm{d}\varphi \\ &\geq -\int_{\hat{\theta}}^{\theta} K_2[\overline{y} - y^*(\varphi), \varphi]\mathrm{d}\varphi + \int_{\hat{\theta}}^{\theta} \underbrace{\{K_1[\overline{y} - y^*(\varphi), \varphi] - \tau_1\}}_{=0} \frac{\mathrm{d}y^*(\varphi)}{\varphi}\mathrm{d}\varphi,\end{aligned}$$

establishing global incentive compatibility.

University of Bonn
University of Frankfurt and Imperial College London
Frankfurt School of Finance and Management

Accepted: 31 December 2015

Additional Supporting Information may be found in the online version of this article:

Appendix B. Non-linear Tax with Loan-based Grants.
Appendix C. Taxes on Output.

References

Akerlof, G. (1978). 'The economics of "tagging" as applied to the optimal income tax, welfare programs, and manpower planning', *American Economic Review*, vol. 68(1), pp. 8–19.
Arnott, R. and Stiglitz, J.E. (1988). 'Randomization with asymmetric information', *Rand Journal of Economics*, vol. 19(3), pp. 344–62.
Boadway, R. and Tremblay, J.F. (2004). 'Public economics and start-up entrepreneurship', in (V. Kanniainen and C. Keuschnigg, eds.), *Venture Capital, Entrepreneurship and Public Policy*, pp. 198–219, Cambridge, MA: MIT Press.
Bovenberg, A.L. and de Mooij R.A. (1994). 'Environmental levies and distortionary taxation', *American Economic Review*, vol. 84(4), pp. 1085–9.

Boyer, M., Hiriart, Y. and Martimort, D. (2006). *Frontiers in the Economics of Environmental Regulation and Liability*, Aldershot and Burlington: Ashgate Publishing.

Boyer, M. and Laffont, J.-J. (1997). 'Environmental risks and bank liability', *European Economic Review*, vol. 41 (8), pp. 1427–59.

Cremer, H., Gahvari, F. and Ladoux, N. (1998). 'Externalities and optimal taxation', *Journal of Public Economics*, vol. 70(3), pp. 343–64.

Diamond, P.A. (1973). 'Consumption externalities and imperfect corrective pricing', *Bell Journal of Economics and Management Science*, vol. 4(2), pp. 526–38.

Diamond, P.A. and Mirrlees, J.A. (1971). 'Optimal taxation and public production II: tax rules', *American Economic Review*, vol. 61(3), pp. 261–78.

Goulder, L.H. and Bovenberg A.L. (2002). 'Environmental taxation and regulation in a second-best setting', in (M. Feldstein and A.J. Auerbach, eds.), *Handbook of Public Economics* vol. 3, pp. 1471–545, Amsterdam: Elsevier.

Hiriart, Y. and Martimort, D. (2006). 'The benefits of extended liability', *Rand Journal of Economics*, vol. 37(3), pp. 562–82.

Jacobs, B. and de Mooij, R.A. (2011). 'Pigou meets Mirrlees: on the irrelevance of tax distortions for the second-best Pigouvian tax', CES Ifo Working Paper 3342.

Kaplow, L. (2012). 'Optimal control of externalities in the presence of income taxation', *International Economic Review*, vol. 53(2), pp. 487–509.

Kopczuk, W. (2003). 'A note on optimal taxation in the presence of externalities', *Economics Letters*, vol. 80(1), pp. 81–6.

Mankiw, N.G. and Weinzierl, M. (2010). 'The optimal taxation of height: a case study of utilitarian income redistribution', *American Economic Journal: Economic Policy*, vol. 2(1), pp. 155–76.

Mirrlees, J.A. (1971). 'An exploration in the theory of optimum income taxation', *Review of Economic Studies*, vol. 38(2), pp. 175–208.

Pigou, A. (1920). *The Economics of Welfare*, London: Macmillan.

Pitchford, R. (1995). 'How liable should a lender be? The case of judgement-proof firms and environmental risk', *American Economic Review*, vol. 85(5), pp. 1171–186.

Rothschild, C. and Scheuer, F. (2014*a*). 'A theory of income taxation under multidimensional skill heterogeneity', Working Paper, Wellesley College and Stanford University.

Rothschild, C. and Scheuer, F. (2014*b*). 'Optimal taxation with rent-seeking', Working Paper, Wellesley College and Stanford University.

Sandmo, A. (1975). 'Optimal taxation in the presence of externalities', *Swedish Journal of Economics*, vol. 77(1), pp. 86–98.

Seade, J.K. (1977). 'On the shape of optimal tax schedules', *Journal of Public Economics*, vol. 7(2), pp. 203–35.

Seade, J.K. (1982). 'On the sign of the optimum marginal income tax', *Review of Economic Studies*, vol. 49(4), pp. 637–43.

Tirole, J. (2010). 'From Pigou to extended liability: on the optimal tax of externalities under imperfect financial markets', *Review of Economic Studies*, vol. 77(2), pp. 697–729.