TI 2024-035/VIII
Tinbergen Institute Discussion Paper

Towards fully decentralized environmental regulation

Jens Gudmundsson[1]
Jens Leth Hougaard[2]
Erik Ansink[3]

1 University of Copenhagen

2 University of Copenhagen

3 Vrije Universiteit Amsterdam and Tinbergen Institute

Towards fully decentralized environmental regulation

Decentralized environmental regulation

Jens Gudmundsson[1,*], Jens Leth Hougaard[1], and Erik Ansink[2,3]

Abstract: We take a decentralized approach to regulating environmental pollution in settings where each agent's pollution possibly affects all others. There is no central agency to enforce pollution abatement or coordinate monetary transfers. Moreover, agents possess private information, which precludes deducing efficient abatement in general. We propose to implement transfer schemes through smart contracts to allow beneficiaries to compensate for abatement. We characterize all schemes that induce efficient abatement in unique dominant-strategy equilibrium. Moreover, appealing to classical fairness tenets, we pin down the "beneficiaries-compensates principle". Supporting this principle through smart contracts provides a promising step towards decentralized coordination on environmental issues.

Keywords: Pollution; Decentralization; Smart contracts; Beneficiaries-compensates principle

Classification: C72; D62; Q52; H23

1 Introduction

Traditional economic solutions to environmental protection often depend on regulatory agencies to set standards or to implement economic instruments such as emission taxes. However, as environmental issues transcend national borders and jurisdictions, pinpointing a single overseeing agency is often problematic. Without this agency, strategies in the Pigouvian tradition such as the widely-accepted "polluter-pays" principle (Ambec and Ehlers, 2016) cannot really work—there is no authority that can penalize polluters. And even with such agencies, they might lack necessary data to make the right decisions. This rules out, for instance, relying on partially-informed authorities to offer participants compensation to disclose vital, private information (Montero, 2008) needed for efficient decision-making. Suffice to say, the

---

[*] **Correspondence address:** Rolighedsvej 23, 1958 Frederiksberg C, Denmark; jg@ifro.ku.dk.
This version: May 22, 2024.

effectiveness of a centralized authority is debated (see e.g. Sigman, 2005; D'Haultfoeuille *et al.*, 2014; Huse and Lucinda, 2014; Cai *et al.*, 2016). In response, this paper proposes a decentralized approach to the problem of regulating environmental pollution.

We introduce a stylized model in which pollution abatement by one can affect many. This captures a broad range of environmental settings ranging from "global" pollution problems (e.g., greenhouse gases) to ones with a more intricate relational structure (e.g., water pollution in river deltas). Naturally, agents may pollute and be victims of pollution at the same time. We envision a setting where there is no central agency to enforce pollution abatement. Absent such a regulatory agency (or any mutual agreement to do otherwise), a realistic default position is that each party acts in its own best interest. Therefore, costly abatement will typically only be practically feasible if it is financially supported for by those who benefit from it. In this way, solutions where victims compensate polluters may often be the only viable option for international environmental improvements (Buchholz and Rübbelke, 2019). Hence, contrary to the centralized setting, a decentralized setting calls for some version of the "victim-pays" principle (or "pollutee-pays" principle, see e.g. Huber and Wirl, 1998).

Financial support for abatement comes in the form of monetary transfers between agents. We will argue that such transfers are viable also in a decentralized setting and, in particular, suggest to implement them using blockchain-based *smart contracts*. Our aim is to systematically select transfers that always incentivize efficient abatement. This is further complicated by informational asymmetry: an agent's abatement cost is private information, so the agent's efficient abatement level cannot be deduced by anyone except the agent herself. Our first result, Proposition 1, characterizes all transfer schemes that ensure efficient implementation. Formally, the property *efficient implementation* requires that efficient abatement is always the unique dominant-strategy equilibrium in the induced strategic abatement game.

Still, this class of efficient transfer schemes is large, which allows us to impose additional selection criteria. We first normalize transfers to the point that, if there is no abatement to compensate for, then consequently there should be no compensatory transfers either. This property, *no transfers without abatement*, captures the natural status quo of our decentralized setting. Intuitively, we cannot require non-zero transfers for zero abatement as no net-payer would agree to it. In addition, we impose two standard properties that pertain to the stability and consistency of the solution as such, *additivity* and *zero truncation*. These have to do with the transfer scheme being invariant to seemingly arbitrary ways of delimiting the problem (Shapley, 1953; Moulin, 2002). Our main result, Theorem 1, pins down what we denote the "beneficiary-compensates" (BC) principle as the only efficient transfer scheme to satisfy these three desirable axioms. The BC principle is a specific interpretation of the victim-pays principle under which abatement compensation matches the benefits the abatement generates.

We acknowledge that the victim-pays principle has received extensive criticism on moral and ethical grounds, mostly so in the literature that discusses the applicability of the Coase theorem to environmental problems (see Medema, 2014, 2020, for extensive reviews). Randall (1974) famously commented on the possibility that victims would pay polluters as the *'amoralization of the externality issue'*. None of the criticism, however, has been targeted at its economic performance. Indeed, this is corroborated by our results that highlight the efficiency of the BC principle. We note also that, in practice, environmental applications of the Coase theorem often are close in spirit to the victim-pays principle. Examples include projects on payments-for-ecosystem services (Wunder and Albán, 2008), permit trading schemes (Kruger *et al.*, 2007), and negotiated international environmental agreements (Barrett, 2003; Libecap, 2014). A comprehensive overview of applications is given by Deryugina *et al.* (2021).

We propose to use smart contracts to facilitate Coasean-style bargaining. Smart contracts are digital protocols that automatically execute, control, and document events and actions on a blockchain based on predefined conditions (e.g. Saleh, 2021; Cong and He, 2019; Huberman *et al.*, 2021; Halaburda *et al.*, 2022). Concretely, it is a piece of code that governs a set of variables and provides functions to modify these variables. The code is publicly available and can be inspected by all parties before use to ensure that it works as intended. Interactions with the contract occur through *transactions*, which may specify functions (in the contract) to run as well as inputs to run them on. An elementary feature is that a transaction may transfer value between accounts through an associated cryptocurrency. This can for instance be from the user to the contract (say as a deposit) or the other way around (say by calling a "refund" function within the contract that returns the deposit from the contract's account). For efficiency purposes, transactions are grouped together and ran sequentially in *blocks*. The blocks are then chained cryptographically to form the *blockchain*. The variables of the contract are initialized when the contract is "deployed" on the blockchain; the contract then obtains a unique address and its code is fixed. In this way, users are safe in knowing that no one can "override" the contract and make it do something beyond its intended functionalities—no one can for instance empty the contract's balance unless there is a function specifically for this purpose. In our proposal, the contract is quite simple: agents commit to abatement levels, deposit funds to cover potential transfers, and finally receive their allotted transfer once they provide evidence of abatement (we postpone the practical details on this point to Section 4). Still, there are technical issues to overcome, such as how to carry out simultaneous-move games in the inherently "sequential" blockchain setting where earlier transactions are available to later agents.

Smart contracts relieve obstacles related to commitment, incomplete information, and transaction costs (Lewis, 1996; Huber and Wirl, 1998; Cohen and Santhakumar, 2007). A series of pilot projects reviewed by Kotsialou *et al.* (2022) supports this argument by illustrating that smart contracts are increasingly gaining ground in the environmental domain. In short, we view these contracts as ideal decentralized replacements for the central authority prevalent in traditional environmental regulation. Nevertheless, some challenges remain that mostly relate to strategic behaviour. One of those is extortion, where polluters may pollute more to receive more compensation for abatement. Another pertains to free-riding. As agents cannot be forced to participate, they may choose to withdraw from a potential agreement altogether. This stems from the expectation that others will still continue to participate, allowing the withdrawing agent to reap the benefits of the resulting "partial" agreement while evading any corresponding costs or obligations. In this sense, smart contracts cannot solve the self-enforcement problem (Barrett, 1994) that is inherent to essentially all externality settings.

The remainder of this paper is structured as follows. In Section 2, we introduce the model and the BC principle. In Section 3, we first examine efficient implementation (Proposition 1) and then explore desirable normative conditions to characterize the BC principle (Theorem 1). In Section 4, we propose a practical implementation through smart contracts. We conclude in Section 5.

2 Preliminaries

We explore a setting in which a group of agents (countries, regions, citizens across jurisdictions) are engaged in pollution-generating activities. Abatement by one may benefit many, so beneficiaries may be willing, but cannot agree on how, to compensate those who abate. To further complicate matters, there is no central authority to enforce any agreement.

2.1 Model

There is a finite set of **agents** N. Throughout, we reserve i and j to denote generic agents. Each agent i can abate $x_i \geq 0$ units of pollution; let $x = (x_i)_{i \in N} \in \mathbb{R}^N_{\geq 0} \equiv X$ be an abatement profile, a **profile** for short. The cost to agent i of abating $x_i \geq 0$ is $\hat{C}_i(x_i) \geq 0$ and this is known only to agent i. (The hat notation is used throughout to emphasize private information and thus aspects of the problem that cannot be part of any form of "solution".) Each **cost function** $\hat{C}_i \colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is strictly convex and differentiable with zero costs and marginal costs at zero: $\hat{C}_i(0) = \hat{C}_i'(0) = 0$. Cost functions may differ between agents and are fixed throughout. Analogous to Ambec and Ehlers (2016), we assume constant marginal **benefit** $b_{ij} \geq 0$ to agent j of agent i's abatement. That is to say, if i abates x_i units, then agent j enjoys (monetary) benefit $b_{ij}x_i$. We isolate the cost of abatement from its benefit in that agent i incurs the cost but agents $j \neq i$ enjoy the benefit. This elucidates that abatement only is undertaken if properly compensated for. (For instance, farmers routinely apply pesticides, unknowingly compromising groundwater quality. However, the immediate repercussions may not impact their own well-being or farm profit.) Formally, we assume $b_{ii} = 0$ for each agent i. The matrix of benefits $B = (b_{ij})_{i,j \in N}$ is common knowledge among agents. Let $\mathcal{B}$ denote the set of benefit matrices:

$$\mathcal{B} = \{B \in \mathbb{R}^{N \times N}_{\geq 0} \mid b_{ii} = 0 \text{ for each } i \in N\}.$$

Monetary transfers are possible and allow beneficiaries to compensate for abatement efforts. For instance, for j to compensate i, agent j may invest in sustainable projects in region i. A transfer scheme t, **transfers** for short, associates to each profile $x \in X$ and benefits $B \in \mathcal{B}$ a vector $t(x, B) \in \mathbb{R}^N$ with $\sum_i t_i(x, B) = 0$. That is, we restrict from the outset to *balanced* transfers. Balance is desirable at it neither requires agents to wastefully "burn" money nor does it make the solution rely on subsidies from external parties. If $t_i(x, B) > 0$, then i receives $t_i(x, B)$; if negative, then i pays $-t_i(x, B)$. We assume each $t_i(x, B)$ is differentiable in x_i. That is, small abatement changes only lead to small transfer changes and, intuitively, there are no "kinks" in the transfer schemes.

Benefits, transfers, and costs are measured on a common scale and agent preferences are quasilinear. At profile x, benefits B, and transfers t, agent i's (privately known) **payoff** is

$$\hat{u}_i(x, B, t) = \textstyle\sum_j b_{ji} x_j + t_i(x, B) - \hat{C}_i(x_i).$$

We define **welfare** as the sum of payoffs, where the balanced transfers t vanish:

$$\hat{W}(x, B) = \textstyle\sum_i \hat{u}_i(x, B, t) = \sum_i \sum_j b_{ji} x_j - \sum_i \hat{C}_i(x_i).$$

Our goal is to maximize welfare by incentivizing efficient abatement. Profile $\hat{x} \in X$ is **efficient** if $\hat{W}(\hat{x}, B) \geq \hat{W}(x, B)$ for all profiles $x \in X$. Let $\hat{\mathcal{E}}(B) \subseteq X$ denote efficient profiles given benefits $B \in \mathcal{B}$. The conditions on the cost functions ensure that there is a unique efficient profile $\hat{x}$ and, in particular, that each element $\hat{x}_i$ is determined through the following associated first-order condition:

$$\frac{\partial \hat{W}(\hat{x}, B)}{\partial x_i} = 0 \iff \textstyle\sum_j b_{ij} = \hat{C}'_i(\hat{x}_i). \tag{1}$$

That is, total marginal benefit of i's abatement should match i's marginal cost at $\hat{x}_i$. Hence, agent i, who knows the cost function $\hat{C}_i$ and the commonly known benefits B, can compute their own efficient abatement. No information on other agents' cost functions is needed nor does i have to form any beliefs about the intended abatement of the others. At the same time, $\hat{C}_i$ is not available to any agent $j \neq i$, so $\hat{x}_i$ is known only to i. (This rules out, for instance, any form of coordination among agents $j \neq i$ to "punish" i in the event that i does not abate $\hat{x}_i$.)

Remark 1 (Generalizing benefits). One could imagine a more general formulation in which j enjoys benefit $B_{ij}(x_i)$ of i abating x_i. Our assumption $B_{ij}(x_i) = b_{ij}x_i$ is then a simple but reasonable first-order approximation that may be quite plausible at least within the range of values x_i that are relevant in practice. The key consequence, that i's abatement affects j in a way that is independent of k's actions, can be justified for instance if there are large geographic distances between abatement activities. *End of remark*

2.2 The beneficiaries-compensates principle

Absent a central authority to enforce any abatement whatsoever, a natural proposition is that abatement should be compensated for by its beneficiaries. This proposition is in line with the victim-pays principle. From the set of possible interpretations of this principle, we choose one where compensations equal obtained benefits from pollution abatement. We label this the *beneficiaries-compensates principle*, or BC for short. At profile x, agent i's abatement benefits agent j by $b_{ij}x_i$ whereas j's abatement benefits i by $b_{ji}x_j$; the transfer between the two nets out to $b_{ij}x_i - b_{ji}x_j$. Summing over all agents, this defines transfers t^{BC}:

Definition 1. For each profile $x \in X$, benefits $B \in \mathcal{B}$, and agent i,

$$t_i^{\mathrm{BC}}(x, B) = x_i \textstyle\sum_j b_{ij} - \sum_j b_{ji}x_j.$$

The first term is the total benefit of i's abatement; the second is total benefit that i derives from others' abatement. With these transfers, each agent i's payoff simplifies to the point that it is independent of abatement x_{-i} chosen by the others:[1]

$$\hat{u}_i(x, B, t^{\mathrm{BC}}) = \textstyle\sum_j b_{ji}x_j + x_i \sum_j b_{ij} - \sum_j b_{ji}x_j - \hat{C}_i(x_i) = x_i \sum_j b_{ij} - \hat{C}_i(x_i).$$

Intuitively, t^{BC} is closely related to the "polluter-pays" (PP) principle explored by Ambec and Ehlers (2016). Whereas the PP principle is natural when there is a central authority that can enforce a world without pollution, the BC principle is a complementary solution with particular appeal in the absence of such a central agency. Moreover, we contend that BC intrinsically holds an advantage over PP in terms of incentives to hide or reveal relevant activities. To penalize polluters, the burden of proof lies on the monitoring agency and must

---

[1] As usual, subscript $-i$ projects onto the subspace relative to $N \setminus \{i\}$.

be implemented, for instance, using costly monitoring and enforcement. This is a challenge for many forms of pollution where it is difficult to monitor violations or determine points of emission (Gray and Shimshack, 2011). In contrast, the BC principle shifts burden onto the abating agent who must provide proof of abatement. Under PP, agents are best off polluting out of sight in the middle of the night; under BC, all abatement takes place in broad daylight for all to see. In this way, going from penalizing polluters to compensating contributors may also be a promising step towards more compliant environmental regulation.

EXAMPLE 1 (Illustrating the concepts). Let $\hat{C}_i(x_i) = \alpha_i x_i^2/2$, so $\hat{C}_i'(x_i) = \alpha_i x_i$ and efficient abatement is $\hat{x}_i = \sum_j b_{ij}/\alpha_i$ by equation 1. For transfers t^{BC}, agent i's payoff at the efficient profile is

$$\hat{u}_i(\hat{x}, B, t^{\mathrm{BC}}) = \hat{x}_i \sum_j b_{ij} - \hat{C}_i(\hat{x}_i) = \frac{1}{\alpha_i}\left(\sum_j b_{ij}\right)^2 - \frac{\alpha_i}{2}\left(\sum_j b_{ij}/\alpha_i\right)^2 = \frac{1}{2\alpha_i}\left(\sum_j b_{ij}\right)^2.$$

Hence, in line with intuition, agents are better off the more cost effective they are (α_i low) and the greater the impact that they have on others (b_{ij} high). *End of example*

3 Fair and efficient implementation

Without a central authority to coordinate abatement, agents take action independently. Given common knowledge on benefits B and transfers t, this defines a non-cooperative game $\Gamma(B,t)$ played by agents N in which all simultaneously choose their own abatement level. As outlined in Section 2, each agent i selects abatement $x_i \geq 0$ and obtains payoff $\hat{u}_i(x, B, t)$ at profile x:[2]

$$\hat{u}_i(x, B, t) = \sum_j b_{ji} x_j + t_i(x, B) - \hat{C}_i(x_i).$$

A **dominant-strategy equilibrium** $x^* \in X$ of the game $\Gamma(B,t)$ is such that, for each agent i and profile $x \in X$,

$$\hat{u}_i((x_i^*, x_{-i}), B, t) \geq \hat{u}_i(x, B, t).$$

That is, abating x_i^* is dominant for each agent i. In terms of predicting strategic behavior, this is a very compelling solution concept: if x_i^* is dominant, then x_i^* is the only abatement level that i can rationalize choosing. Let $\mathcal{D}(B,t) \subseteq X$ be the dominant-strategy equilibria

---

[2] Formally, this is a game of incomplete information where an agent knows only their own payoff. One could introduce a more intricate Bayesian framework with type spaces and priors to implement in Bayesian Nash equilibrium, but we opt for prior-free implementation in dominant strategies. This has been argued to be more robust and poses fewer informational requirement, see e.g. Ledyard (1986).

of $\Gamma(B,t)$. Generally, there need not exist such equilibria, but there are transfers t for which there always are. A simple example is to never impose any non-zero transfers. This defines t^0 with $t^0(x,B) = (0,\dots,0)$ and $\mathcal{D}(B,t) = \{(0,\dots,0)\} \neq \emptyset$.

3.1 Efficient implementation

Our goal throughout is to implement efficient abatement. Specifically, we seek transfers t for which it is always the case that every equilibrium is efficient ("exact implementation"). Again, we stress that each agent i can compute their efficient $\hat{x}_i$—what remains is to design t such that $\hat{x}_i$ is dominant. We say that transfers t satisfy the property of *efficient implementation* whenever the set of dominant-strategy equilibria, $\mathcal{D}(B,t)$, always matches the efficient profiles, $\hat{\mathcal{E}}(B)$.

Axiom 1 (Efficient implementation). For each $B \in \mathcal{B}$,

$$\mathcal{D}(B,t) = \hat{\mathcal{E}}(B).$$

Our first result provides a characterization of such transfers. In particular, Proposition 1 pins down how agent i's transfer $t_i(x,B)$ must depend on i's own abatement and, furthermore, how it must separate i's abatement x_i from the abatement of the others, x_{-i}.

Proposition 1. *Transfers t satisfy efficient implementation if and only if there is $f\colon X \times \mathcal{B} \to \mathbb{R}^N$ such that, for each profile $x \in X$, benefits $B \in \mathcal{B}$, and agent i,*

$$t_i(x,B) = x_i \textstyle\sum_j b_{ij} + f_i(x,B),$$

where $f_i(x,B)$ is independent of x_i and $\sum_i f_i(x,B) = -\sum_i x_i \sum_j b_{ij}$.

Proof. Let $x \in X$, $B \in \mathcal{B}$, and consider transfers t. As agent i's payoff is $\hat{u}_i(x,B,t) = \sum_j b_{ji}x_j + t_i(x,B) - \hat{C}_i(x_i)$, i's individual first-order condition is

$$\frac{\partial \hat{u}_i(x,B,t)}{\partial x_i} = 0 \iff \frac{\partial t_i(x,B)}{\partial x_i} = \hat{C}_i'(x_i).$$

Recall also the first-order condition for efficient abatement $\hat{x}$ (equation 1):

$$\frac{\partial \hat{W}(\hat{x},B)}{\partial x_i} = 0 \iff \textstyle\sum_j b_{ij} = \hat{C}_i'(\hat{x}_i).$$

Hence, abating $\hat{x}_i$ is dominant for i if and only if $\partial t_i(x,B)/\partial x_i = \sum_j b_{ij}$. Therefore, we must have $t_i(x,B) = x_i \sum_j b_{ij} + f_i(x,B)$ with f_i independent of x_i. By balance of t, f is such that $\sum_i f_i(x,B) = -\sum_i x_i \sum_j b_{ij}$. □

For transfers t that satisfy *efficient implementation* with associated function f, agent i's payoff is

$$\begin{aligned}\hat{u}_i(x,B,t) &= \textstyle\sum_j b_{ji}x_j + t_i(x,B) - \hat{C}_i(x_i)\\ &= \textstyle\sum_j b_{ji}x_j + x_i\sum_j b_{ij} + f_i(x,B) - \hat{C}_i(x_i).\end{aligned}$$

Most transfer schemes fail to incentivize efficient abatement (e.g., t^0 defined above), but there are also some that succeed in this. For instance, transfers t^{BC} satisfy the condition of Proposition 1 with the associated function f such that $f_i(x,B) = -\sum_j b_{ji}x_j$ (which indeed is independent of x_i as $b_{ii}=0$). Hence, t^{BC} always induces games in which abating efficiently is a dominant strategy for everyone.

3.2 Fairness concerns

Still, there are many transfers beyond this that satisfy *efficient implementation.* Suppose, for instance, we define transfers t^1 equal to t^{BC} except that we pick two agents, say we label them 1 and 2, and everywhere set $t^1_1(x,B) = t^{\mathrm{BC}}_1(x,B)+1$ and $t^1_2(x,B) = t^{\mathrm{BC}}_2(x,B)-1$. As the fixed additional unit transfer has no effect on incentives, t^1 also meets our requirement of implementing efficient abatement. However, this fails the rudimentary principle that if there is no abatement, then there also should be no transfers:

Axiom 2 (No transfers without abatement). For each $B \in \mathcal{B}$,

$$t((0,\dots,0),B) = (0,\dots,0).$$

In contrast, it is immediate that t^{BC} satisfies *no transfers without abatement.*

There are yet more arguments in favor of t^{BC}. A fundamental requirement is that only agents truly pertinent to the problem at hand should matter. For instance, whether an agent who neither benefits from abatement nor harms others is included in the agreement or not should not matter for the others' transfers. If this property is not guaranteed, then there is room for disputes on how to define the boundaries of the problem, which could jeopardize the stability of the agreement. Our *zero truncation* is a variation of the "dummy", independence, and consistency principles that form cornerstones of the literature on fair allocation (Shapley, 1953; Arrow, 1963; Moulin, 2004; Thomson, 2019; Gudmundsson *et al.*, 2023).

Formally, for benefits $B \in \mathcal{B}$ and agent i, we let B_{-i} denote the $(n-1)\times(n-1)$ matrix obtained by eliminating row i and column i from B.[3] That is, we are left with the entries

---

[3] Technically, we now extend to a variable-population model (see e.g. Thomson, 1990) with an infinite set of "potential" agents indexed by the natural numbers $\mathbb{N}$. To specify a problem, we first draw a finite number of

b_{jk} for $j \neq i$ and $k \neq i$. This operation is only straightforward to analyze if all eliminated elements are zero, which is precisely when *zero truncation* applies. As transfers are balanced, *zero truncation* implies that, for an agent i with $b_{ij} = b_{ji} = 0$ for all agents j, we have $t_i(x, B) = 0$.

Axiom 3 (Zero truncation). For each $x \in X$, $B \in \mathcal{B}$, and $i \in N$,

$$\big(b_{ij} = b_{ji} = 0 \text{ for each } j \in N\big) \implies \big(t_j(x, B) = t_j(x_{-i}, B_{-i}) \text{ for each } j \neq i\big)\,.$$

Zero truncation and our final property, *additivity in benefits*, are both satisfied by t^{BC}. Additivity is another staple in the literature on fair allocation, tracing back to Shapley (1953) and explored extensively in axiomatic work ever since (e.g., Moulin, 2002; Bergantiños and Moreno-Ternero, 2020; Gudmundsson *et al.*, 2023). In essence, if abatement has multiple benefits, it should make no difference if we account for them in separate or all at once (e.g., water pollution abatement can improve both water quality and biodiversity). Again, if this property is not met, it opens for undesired conflicts on how to delimit the problem.

Axiom 4 (Additivity in benefits). For each $x \in X$ and $B, B' \in \mathcal{B}$,

$$t(x, B) + t(x, B') = t(x, B + B').$$

We are now ready to present our main result. It is easy to see that t^{BC} satisfies all four requirements but, even more, Theorem 1 shows that it is the *only* transfer scheme to do so. That is to say, if the desire is to implement efficient abatement in a way that only imposes compensatory transfers if there is abatement to compensate for and that is robust to seemingly arbitrary disagreements on how to set the boundaries of the problem, then the only scheme that will do is t^{BC}.

Theorem 1. *Transfers t satisfy efficient implementation, no transfers without abatement, zero truncation, and additivity in benefits if and only if $t = t^{BC}$.*

Proof. As argued, it is immediate that transfers t^{BC} satisfy the properties. We turn therefore to the statement's other direction.

Let transfers t satisfy the properties of the statement. We will show that $t = t^{\mathrm{BC}}$. Let $x \in X$ and $B \in \mathcal{B}$. For each pair of agents i and j, define the $n \times n$ matrix of benefits

---

them from this infinite population. Let $\mathcal{N}$ denote the family of nonempty finite subsets of $\mathbb{N}$. Then $N \in \mathcal{N}$ is a generic set of agents. Let further $X^N = \mathbb{R}^N_{\geq 0}$ and $\mathcal{B}^N = \mathbb{R}^{N \times N}_{\geq 0}$ denote the sets of profiles and benefits that agents $N \in \mathcal{N}$ may face. A transfer scheme t associates to each population $N \in \mathcal{N}$, profile $x \in X^N$, and benefit matrix $B \in \mathcal{B}^N$ a balanced vector of transfers $t(x, B) \in \mathbb{R}^N$. Our axioms should be understood as applying to every population $N \in \mathcal{N}$.

$B^{ij} \in \mathcal{B}$ in which all entries are zero except b_{ij}. (Of course, if also $b_{ij} = 0$, then B^{ij} only contains zeros.) Intuitively, the condition for *zero truncation* applies for each agent $k \neq i, j$ in B^{ij}. Eliminating all of the rows and columns associated to agents $k \neq i, j$, we obtain the two-by-two matrix $\bar{B}^{ij}$. This procedure is illustrated below for $n = 4$ and $i, j = 2, 3$.

$$B = \begin{bmatrix} 0 & b_{12} & b_{13} & b_{14} \\ b_{21} & 0 & b_{23} & b_{24} \\ b_{31} & b_{32} & 0 & b_{34} \\ b_{41} & b_{42} & b_{43} & 0 \end{bmatrix} \qquad B^{23} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & b_{23} & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \longleftarrow \bar{B}^{23}$$

By repeatedly applying *zero truncation*, we have $t_i(x, B^{ij}) = t_i((x_i, x_j), \bar{B}^{ij})$ and $t_j(x, B^{ij}) = t_j((x_i, x_j), \bar{B}^{ij})$, whereas $t_k(x, B^{ij}) = 0$ for each $k \neq i, j$. By *efficient implementation* through Proposition 1, there is a function f such that

$$t_i((x_i, x_j), \bar{B}^{ij}) = b_{ij}x_i + f_i((x_i, x_j), \bar{B}^{ij}) \quad \text{and} \quad t_j((x_i, x_j), \bar{B}^{ij}) = f_j((x_i, x_j), \bar{B}^{ij}),$$

where f_i and f_j are independent of x_i and x_j, respectively. By balance, $f_i((x_i, x_j), \bar{B}^{ij}) + f_j((x_i, x_j), \bar{B}^{ij}) = -b_{ij}x_i$. But then f_i cannot depend on x_j as otherwise f_j would as well, contradicting the design of f. Hence, f_i is independent of both x_i and x_j, so in particular $f_i((x_i, x_j), \bar{B}^{ij}) = f_i((0, 0), \bar{B}^{ij})$. By *no transfers without abatement*, $f_i((0, 0), \bar{B}^{ij}) = 0$. Hence, $t_i((x_i, x_j), \bar{B}^{ij}) = -t_j((x_i, x_j), \bar{B}^{ij}) = b_{ij}x_i$. In summary,

$$t_\ell(x, B^{ij}) = \begin{cases} b_{ij}x_i & \text{for } \ell = i \\ -b_{ij}x_i & \text{for } \ell = j \\ 0 & \text{for } \ell \neq i, j. \end{cases}$$

By construction, $B = \sum_i \sum_j B^{ij}$. By repeatedly applying *additivity in benefits*, for each agent j,

$$\begin{aligned} t_j(x, B) &= \textstyle\sum_i \sum_k t_j(x, B^{ik}) = \sum_i t_j((x_i, x_j), \bar{B}^{ji}) + \sum_i t_j((x_i, x_j), \bar{B}^{ij}) \\ &= x_j \textstyle\sum_i b_{ji} - \sum_i x_i b_{ij} = t_j^{\mathrm{BC}}(x, B). \end{aligned} \qquad \square$$

Remark 2 (Logical independence). There are transfers different from t^{BC} that satisfy any combination of three of the four axioms. Transfers t^0 (always zero everywhere) fail *efficient implementation* but otherwise satisfy the axioms of Theorem 1. We design t^2 and t^3 from t^{BC} by adding a transfer between agents 1 and 2 that is independent of x_1 and x_2 (to ensure that incentives are not distorted and we maintain *efficient implementation*). For convenience, we repeat t^1 below, which failed *no transfers without abatement.*

$$\begin{aligned} t_1^1(x,B) &= t_1^{\mathrm{BC}}(x,B) + 1 \\ t_1^2(x,B) &= t_1^{\mathrm{BC}}(x,B) + x_n \\ t_1^3(x,B) &= t_1^{\mathrm{BC}}(x,B) + b_{1n}^2 x_n \end{aligned}$$

For each $\ell \in \{1,2,3\}$, we set $t_2^\ell(x,B) = t_2^{\mathrm{BC}}(x,B) - (t_1^\ell(x,B) - t_1^{\mathrm{BC}}(x,B))$ and let $t_k^\ell(x,B) = t_k^{\mathrm{BC}}(x,B)$ for all agents $k \neq 1,2$.

Transfers t^2 are "bossy" (Satterthwaite and Sonnenschein, 1981; Svensson, 1999) in that agent n can affect agent 1 and 2's transfers without affecting her own. These transfers fail *zero truncation*: if $b_{jn} = b_{nj} = 0$ for all agents j—so agent n is "irrelevant"—then $t_n^2(x,B) = 0$ yet $t_1^2(x,B)$ and $t_2^2(x,B)$ depend on x_n. For t^3, we recover *zero truncation* by multiplying by b_{1n} but, as the term is squared, t^3 fails *additivity in benefits.* *End of remark*

As already noted, a property of t^{BC} is that an agent's payoff is independent of the abatement choices of the others. Such payoff independence may hold little normative appeal in itself, but it can be desirable in practice. It leaves each agent in full control of their own well-being and removes any need to form beliefs about others. Moreover, it implies that no group of agents can collude to manipulate the system as each agent only affects their own payoff.[4] Next, we sketch how t^{BC} is the only transfer scheme with this independency that also satisfies *efficient implementation* and *no transfers without abatement.*

Recall that transfers t satisfying *efficient implementation* with associated function f imply that agent i's payoff is

$$\hat{u}_i(x,B,t) = \textstyle\sum_j b_{ji}x_j + x_i \sum_j b_{ij} + f_i(x,B) - \hat{C}_i(x_i).$$

We now ask more generally what has to hold for this to be independent of x_j for each $j \neq i$. Clearly, the first term $\sum_j b_{ji}x_j$ has to cancel—there has to exist a function $g\colon \mathcal{B} \to \mathbb{R}^N$ such that $f_i(x,B) = g_i(B) - \sum_j b_{ji}x_j$.[5] Then

$$\hat{u}_i(x,B,t) = x_i \textstyle\sum_j b_{ij} + g_i(B) - \hat{C}_i(x_i) = \hat{u}_i(x,B,t^{\mathrm{BC}}) + g_i(B).$$

---

[4] For instance, had agent j's payoff increased more than agent i's decreased when i abated inefficiently, then j could promise side-transfers to i that would make both better off, undermining efficient implementation.

[5] To achieve the desired payoff independence, g_i cannot depend on x_j for $j \neq i$. Moreover, as f_i does not depend on x_i, also g_i does not. Thus, g_i is independent of x.

If we now impose *no transfer without abatement*, then $\hat{u}_i((0,\ldots,0),B,t) = g_i(B) = 0$ for each $B \in \mathcal{B}$ and, in consequence, $t = t^{\mathrm{BC}}$. Hence, this offers an alternative characterization in which payoff independence replaces Theorem 1's *zero truncation* and *additivity in benefits*.

4 Decentralized coordination

As shown above, we can theoretically ensure efficient implementation using the BC principle, which also stands out as a natural alternative in the absence of a trusted third party. However, there are clear commitment issues to overcome in practice: if agent i should abate to agent j's benefit, then

- Agent i prefers to be compensated upfront, as there is no one to penalize agent j if j simply chooses not to pay afterwards;
- Agent j prefers to pay afterwards, as there is no one to penalize agent i if i simply chooses to take the upfront payment and not abate.

In what follows, we propose a solution to this issue—credible commitment in a decentralized context—using blockchain technology and its smart contracts. The idea is simple: the smart contract acts as a trusted middleman that takes j's deposit upfront and, depending on i's actions, either compensates i for abatement or returns the funds to j if i fails to honor their promise. This is cryptographically fixed in code and runs in an automated, self-executing way to alleviate any need for mutual trust between i and j. This confirms what was noted by Gans (2021), namely that inherent trust-promoting features of blockchains can counteract real-world trust deficits. Moreover, it also has potential to reduce transaction costs and increase transparency (see also Bakos and Halaburda, 2022; Brzustowski *et al.*, 2023).

4.1 Smart-contract implementation of BC

Using a smart contract to implement t^{BC} will involve a procedure in several steps. We suggest a simple such approach in which agents first choose abatement levels; beneficiaries then deposit enough to compensate abating agents; and finally, abating agents who provide timely evidence of abatement (e.g., corroborating sensor data) receive compensation through the deposits. We detail all of this in Definition 2 below, but before that we elaborate on a technical challenge associated to the procedure's first step.

A difficulty in moving a strategic game onto the inherently sequential blockchain is in the simultaneous choice of abatement. Agents interacting later with the contract inevitably will observe earlier agents' interactions. Hence, it is crucial that there is a way for an agent to commit to abating x_i without anyone else being able to infer x_i. This is resolved by dividing the step in two—an encryption and a decryption part—linked through a *cryptographic hash function*.[6] Thus, say agent i wants to abate $x_i \geq 0$. If the agent announces x_i outright, then this would be observable to all others who could condition their choice on x_i and we would not capture the desired strategic structure. Therefore, we take two measures to conceal x_i using a publicly known hash function H and a parameter $M \geq 0$. Specifically, the agent first chooses some $x_i' \geq 0$ such that $x_i' \bmod M = x_i$ and submits the hash value $H(x_i')$ to the contract. (For instance, with $M = 10$, $x_i' = 27$ corresponds to $x_i = (27 \bmod 10) = 7$.) This accomplishes two things. Other agents cannot infer x_i (or x_i') from the encrypted commitment. Moreover, it provides many ways for i to commit to x_i.[7] Once all encrypted commitments have been submitted, i submits x_i' to the contract, which verifies that the hash value of this submission matches the previous encrypted commitment. In summary, driven by the strong security properties of cryptographic hash functions as elaborated in footnote 6,

- In the encryption step, even though agent j may have observed $H(x_i')$ when choosing j's own encrypted commitment, j cannot infer x_i and condition its choice on it;
- In the decryption step, even though agent j may have learned i's chosen abatement x_i before revealing x_j', j is unable to find a different input y_j' to match its encrypted commitment (that is, with $H(y_j') = H(x_j')$).

Hence, even though agents inevitably act sequentially, the information structure is equivalent to that of a simultaneous-move game. As M is publicly known, the profile x of intended abatement is known to all agents once x' has been submitted. Next, we present a bare-bones version of the contract.

Definition 2 (Smart-contract implementation of BC). Benefits $B \in \mathcal{B}$, hash function H, and parameter $M \geq 0$ are public information. The contract operates in the following steps:

1. Agents independently decide on abatement levels. Agent i chooses abatement $x_i \geq 0$ as well as some artificial $x_i' \geq 0$ such that $x_i' \bmod M = x_i$ and submits the encrypted $H(x_i')$ to the contract's "encrypted commitment" function. Once all encrypted commitments have been submitted, we proceed to the next step.

---

[6] This is a deterministic function H that maps arbitrary inputs to outputs of a fixed size. Inputs are mapped "evenly" over the output range and small input variations lead to "unpredictable" output changes: a sequence of inputs $a_1, a_2, \ldots$ produce outputs $H(a_1), H(a_2), \ldots$ that appear drawn uniformly from the output range. Computing $H(x)$ is easy, but reverse-engineering an input x from the encrypted output $H(x)$ or finding a different input y with the same output, $H(x) = H(y)$, is intractable.

[7] If agents submit $H(x_i)$ directly and there are reasons to expect x_i to be bounded by $K \geq 0$, then it suffices to compute $H(z)$ for $z \leq K$ to learn x_i. This is done quickly unless K is very large. In contrast, our approach gives many more ways of "representing" x_i through x_i', so "checking all of them" is no longer viable.

2. Each agent i submits x'_i to the "decrypt commitment" function. The contract verifies that the encrypted commitment matches the hash of the actual commitment. Moreover, the contract computes $x_i = x'_i \bmod M$. Once all actual commitments have been submitted, we proceed to the next step.
3. As the profile x now is public and benefits B are common knowledge, the contract can compute how much each agent should compensate others. Each agent i now transfers (at least) this amount $\sum_j x_j b_{ji}$ to the contract.
4. When agent i submits proof of abatement, the amount $x_i \sum_j b_{ij}$ is transferred from the contract to i. If instead agent i fails to submit proof of abatement in time, then, for each agent j, the amount $x_i b_{ij}$ is transferred from the contract to j.

In practice, one may want to extend the contract with aspects such as initial deposits (distributed among the others if some agent misbehaves) and deadlines at which the contract reverts and returns deposits. We remark also that, with minor modifications, Definition 2 allows implementation of other transfer schemes beyond t^{BC}.

A prerequisite for this approach to work is that benefits B are hard-coded in the contract. An argument in favor of this is that B is common knowledge and that the contract is open-source and easy to construct. Hence, any agent can put together the contract, and if there is disagreement on its contents, any agent can easily propose a new one. Only once all have agreed on the terms (so B is fixed and available) do we proceed with the above process. Alternatively, some form of "committee-based consensus" (Benhaim *et al.*, 2023) could be used. In essence, each agent i distributes "their" view of B to all others, and all agents j choose whether to approve it (that is, append their cryptographic signature). The matrix B that the contract eventually refers to is the one approved by a large-enough fraction of agents. Further incentive mechanisms could be added such as rewards to approving agents (say, collected as deposits/entrance fees in advance). An interesting question left for future research is whether it is possible to design intuitive, informationally lean mechanisms that incentivize agents to report B truthfully (e.g., each agent reports some subset of B, compare Jackson *et al.*, 1994; Sjöström, 1994).

4.2 Digitizing pollution abatement

A key point that remains to be elaborated on is in our procedure's final step, namely how to provide abatement "evidence" to the digital contract. In an idealized centralized setting, one could expect the central authority to accurately and reliably tackle this. For our case, we propose the next-best solution, namely to instead rely on data collected from digital sensors. (For instance, such sensors are already in place along many rivers to monitor water quality.) Smart contracts can combine with the "Internet of Things" to feed sensor data into the contract (Christidis and Devetsikiotis, 2016; Thomas *et al.*, 2019; Pradhan *et al.*, 2022; Biswas *et al.*, 2023; Bakos and Halaburda, 2023) and automatically execute functions in response. In this way, once an agent i has "proof" of abatement (for instance, sensor data indicating lower levels of pollution), funds can be released from the contract to compensate i. A caveat is that the reported data must be trustworthy and difficult to manipulate: our solution is only as good as the data it is based on. Cryptographic solutions to address this are beyond the scope of the paper, but an obvious first step is to have multiple sensors in largely the same area that all should give agreeing readings.

Our solution has potential to work yet better for software-controlled pollution abatement. Similar to how Tesla's Autopilot software comes installed with the purchase of a car but has to be activated through a monthly subscription, we imagine that abatement equipment is supplied to agents but inoperable until digitally activated. It is straightforward for the equipment manufacturer to use a smart contract for license management; that is, purchasing and activating software licenses can be done on-chain (e.g., once the software is activated, it sends a transaction to register this in the manufacturer's contract). The benefit of this approach is that our abatement contract (detailed below) can interact with the manufacturer's contract. When agent i commits to abating a certain amount, i will also be required to deposit enough to cover the expenses of the corresponding licenses. The abatement contract then uses the deposited funds to purchase licenses through the manufacturer's contract. This makes commitment credible to a degree that would be difficult to attain even with a conventional trusted third party. Figure 1 sketches the timeline.

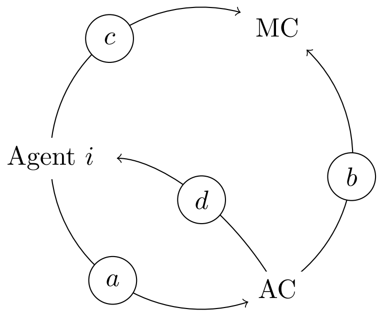

ⓐ Agents interact through the abatement contract (AC) to make commitments and deposits

ⓑ The abatement contract uses the deposits to purchase licenses from the manufacturer (MC)

ⓒ Once the license is activated by the agent, the software registers this in the manufacturer's contract

ⓓ This triggers the abatement contract to pay out the monetary compensation to the agent

Figure 1. *Timeline of agent interaction with the abatement contract and manufacturer's contract.*

Digitization of pollution abatement is already taking place in many domains of environmental management.[8] Ever more accurate and reliable data on pollution and abatement as well as other environmental indicators can be used for environmental regulation. Here, the benefits are clear: digital sensor readings can feed into smart contracts and facilitate more efficient pollution abatement.

5 Concluding remarks

This paper is a step in the direction marked out by Chapron (2017), who asserts that "the environment needs cryptogovernance". Chapron further suggests that the current environmental crisis is fueled in part by a deepening lack of trust. As the gap widens between unfamiliar parties—be it corporations, governments, or consumers—there is a rising risk of policy failure. Such failures are especially likely in the case of regulation by centralized authorities, constrained—besides lack of trust—by incomplete information and transaction costs. We have highlighted one particular solution, the "beneficiaries-compensates" principle, as a promising step towards decentralized environmental protection and elaborated on how smart contracts, a type of cryptogovernance, can take it from theory to practice. This BC principle is attractive in multiple ways. It incentivizes efficient abatement all the while guaranteeing several desirable normative properties. It is particularly suitable and robust in the absence of centralized abatement control and enforcement.

Although the BC principle formally may appear "the opposite" of the polluter-pays principle, we do not view them as necessarily competing regimes. Rather, they present two complementary options with their own advantages depending on the institutional context. The PP principle may from an ethical standpoint seem more appealing—polluters should pay for polluting—whereas the BC principle is more robust. To this point, Buchholz and Rübbelke (2019) argue that "Even though they may seem ethically questionable, solutions where the victim pays a compensation to the polluter may come about faster than those where the polluter has to carry the financial burden of environmental protection. Especially in the international context, environmental improvements frequently become feasible only when the victim pays." Our theoretical analysis suggests that BC is the best one can do in the absence of centralized control in terms of efficiency and fairness. In sum, if it is easy to pinpoint pollution and enforce penalties, then PP may work well; if not, then BC may be a more promising mechanism. Whereas in the former case practical implementation is trivial

---

[8] Several UN Environment Programmes employ data-driven approaches to address diverse environmental challenges. For instance, the *Global Environment Monitoring System for Air* offers real-time estimates of air pollution levels worldwide, aiding in local air quality assessment, the *Freshwater Ecosystems Explorer* provides accessible geospatial data on freshwater ecosystem changes, and the *International Methane Emissions Observatory* delivers reliable data to mitigate methane emissions and combat climate change.

(the centralized agency enforces the solution), we have argued that blockchain-based smart contracts provide an excellent toolkit in the latter case.

Moving forward, the promise of blockchain technology for environmental solutions lies in effectively integrating it with well-thought-out environmental policy design. If successful, the technology has vast potential to "ease the frictions that prevent a vast array of sustainability, humanitarian, and environmental initiatives from fulfilling their potential" (UNDP, 2018).

Affiliations

[1]Department of Food and Resource Economics, University of Copenhagen, Denmark
[2]Department of Spatial Economics, Vrije Universiteit Amsterdam, The Netherlands
[3]Tinbergen Institute, The Netherlands

References

Ambec, S. and Ehlers, L. (2016). 'Regulation via the Polluter-pays Principle', *The Economic Journal*, vol. 126(593), pp. 884–906.

Arrow, K.J. (1963). *Social Choice and Individual Values*, Yale University Press.

Bakos, Y. and Halaburda, H. (2022). 'Overcoming the Coordination Problem in New Marketplaces via Cryptographic Tokens', *Information Systems Research*, vol. 33(4), pp. 1368–1385.

Bakos, Y. and Halaburda, H. (2023). 'Blockchains, Smart Contracts and Connected Sensors: Enforceable Execution vs Better Information', NYU Stern School of Business.

Barrett, S. (1994). 'Self-Enforcing International Environmental Agreements', *Oxford Economic Papers*, vol. 46, pp. 878–894.

Barrett, S. (2003). *Environment and Statecraft: The Strategy of Environmental Treaty-Making*, Oxford: Oxford University Press.

Benhaim, A., Falk, B.H. and Tsoukalas, G. (2023). 'Scaling Blockchains: Can Committee-Based Consensus Help?', *Management Science*, vol. 69(11), pp. 6525–6539.

Bergantiños, G. and Moreno-Ternero, J.D. (2020). 'Sharing the Revenues from Broadcasting Sport Events', *Management Science*, vol. 66(6), pp. 2417–2431.

Biswas, K., Muthukkumarasamy, V., Bai, G. and Chowdhury, M.J.M. (2023). 'A reliable vaccine tracking and monitoring system for health clinics using blockchain', *Nature Scientific Reports*, vol. 13(570).

Brzustowski, T., Georgiadis-Harris, A. and Szentes, B. (2023). 'Smart Contracts and the Coase Conjecture', *American Economic Review*, vol. 113(5), pp. 1334–1359.

Buchholz, W. and Rübbelke, D. (2019). 'Environmental Externalities and Their Internalization Through Voluntary Approaches', in (W. Buchholz and D. Rübbelke, eds.), *Foundations of Environmental Economics*, pp. 13–36, chap. 2, Springer.

Cai, H., Chen, Y. and Gong, Q. (2016). 'Polluting thy neighbor: Unintended consequences of China's pollution reduction mandates', *Journal of Environmental Economics and Management*, vol. 76, pp. 86–104.

Chapron, G. (2017). 'The environment needs cryptogovernance', *Nature*, vol. 545, pp. 403–405.

Christidis, K. and Devetsikiotis, M. (2016). ‘Blockchains and Smart Contracts for the Internet of Things’, *IEEE Access*, vol. 4, pp. 2292–2303.

Cohen, M.A. and Santhakumar, V. (2007). ‘Information disclosure as environmental regulation: A theoretical analysis’, *Environmental and Resource Economics*, vol. 37, pp. 599–620.

Cong, L.W. and He, Z. (2019). ‘Blockchain Disruption and Smart Contracts’, *The Review of Financial Studies*, vol. 32(5), pp. 1754–1797.

Deryugina, T., Moore, F. and Tol, R.S. (2021). ‘Environmental applications of the Coase Theorem’, *Environmental Science & Policy*, vol. 120, pp. 81–88.

D’Haultfoeuille, X., Givord, P. and Boutin, X. (2014). ‘The environmental effect of green taxation: The case of the French bonus/malus’, *The Economic Journal*, vol. 124(578), pp. F444–F480.

Gans, J.S. (2021). ‘The Fine Print in Smart Contracts’, in (M. C. Compagnucci, M. Fenwick and S. Wrbka, eds.), *Smart Contracts Technological, Business and Legal Perspectives*, pp. 41–57, chap. 2, London, UK: Hart Publishing.

Gray, W.B. and Shimshack, J.P. (2011). ‘The effectiveness of environmental monitoring and enforcement: A review of the empirical evidence’, *Review of Environmental Economics and Policy*, vol. 5(1), pp. 3–24.

Gudmundsson, J., Hougaard, J.L. and Ko, C.Y. (2023). ‘Sharing Sequentially Triggered Losses: Automated Conflict Resolution Through Smart Contracts’, *Management Science*, vol. 70(3), pp. 1773–1786.

Halaburda, H., Haeringer, G., Gans, J. and Gandal, N. (2022). ‘The Microeconomics of Cryptocurrencies’, *Journal of Economic Literature*, vol. 60(3), pp. 971–1013.

Huber, C. and Wirl, F. (1998). ‘The polluter pays versus the pollutee pays principle under asymmetric information’, *Journal of Environmental Economics and Management*, vol. 35(1), pp. 69–87.

Huberman, G., Leshno, J.D. and Moallemi, C. (2021). ‘Monopoly without a Monopolist: An Economic Analysis of the Bitcoin Payment System’, *Review of Economic Studies*, vol. 88, pp. 3011–3040.

Huse, C. and Lucinda, C. (2014). ‘The market impact and the cost of environmental policy: Evidence from the Swedish green car rebate’, *The Economic Journal*, vol. 124(578), pp. F393–F419.

Jackson, M.O., Palfrey, T.R. and Srivastava, S. (1994). 'Undominated Nash implementation in bounded mechanisms', *Games and Economic Behavior*, vol. 6(3), pp. 474–501.

Kotsialou, G., Kuralbayeva, K. and Laing, T. (2022). 'Blockchain's potential in forest offsets, the voluntary carbon markets and REDD+', *Environmental Conservation*, vol. 49(3), pp. 137–145.

Kruger, J., Oates, W.E. and Pizer, W.A. (2007). 'Decentralization in the EU emissions trading scheme and lessons for global policy', *Review of Environmental Economics and Policy*, vol. 1(1), pp. 112–133.

Ledyard, J.O. (1986). 'The scope of the hypothesis of Bayesian equilibrium', *Journal of Economic Theory*, vol. 39(1), pp. 59–82.

Lewis, T.R. (1996). 'Protecting the environment when costs and benefits are privately known', *The RAND Journal of Economics*, vol. 27(4), pp. 819–847.

Libecap, G.D. (2014). 'Addressing global environmental externalities: Transaction costs considerations', *Journal of Economic Literature*, vol. 52(2), pp. 424–479.

Medema, S.G. (2014). 'The curious treatment of the Coase theorem in the environmental economics literature, 1960–1979', *Review of Environmental Economics and Policy*, vol. 8(1), pp. 39–57.

Medema, S.G. (2020). 'The Coase theorem at sixty', *Journal of Economic Literature*, vol. 58(4), pp. 1045–1128.

Montero, J.P. (2008). 'A Simple Auction Mechanism for the Optimal Allocation of the Commons', *American Economic Review*, vol. 98(1), pp. 496–518.

Moulin, H. (2002). 'Axiomatic cost and surplus sharing', in (K. J. Arrow, A. K. Sen and K. Suzumura, eds.), *Handbook of Social Choice and Welfare*, pp. 289–357, vol. 1, chap. 6, Elsevier.

Moulin, H. (2004). *Fair division and collective welfare*, MIT Press.

Pradhan, N.R., Singh, A.P., Verma, S., Kavita, Wozniak, M., Shafi, J. and Ijaz, M.F. (2022). 'A blockchain based lightweight peer-to-peer energy trading framework for secured high throughput micro-transactions', *Nature Scientific Reports*, vol. 12(14523).

Randall, A. (1974). 'Coasian externality theory in a policy context', *Natural Resources Journal*, vol. 14, p. 35.

Saleh, F. (2021). ‘Blockchain without Waste: Proof-of-Stake’, *Review of Financial Studies*, vol. 34(3), pp. 1156–1190.

Satterthwaite, M.A. and Sonnenschein, H. (1981). ‘Strategy-Proof Allocation Mechanisms at Differentiable Points’, *The Review of Economic Studies*, vol. 48(4), pp. 587–597.

Shapley, L.S. (1953). ‘A value for n-person games’, *Annals of Mathematics Study*, vol. 2(28), pp. 307–317.

Sigman, H. (2005). ‘Transboundary spillovers and decentralization of environmental policies’, *Journal of Environmental Economics and Management*, vol. 50(1), pp. 82–101.

Sjöström, T. (1994). ‘Implementation in undominated Nash equilibria without integer games’, *Games and Economic Behavior*, vol. 6, pp. 502–551.

Svensson, L.G. (1999). ‘Strategy-proof allocation of indivisible goods’, *Social Choice and Welfare*, vol. 16(4), pp. 557–567.

Thomas, L., Zhou, Y., Long, C., Wu, J. and Jenkins, N. (2019). ‘A general form of smart contract for decentralized energy systems management’, *Nature Energy*, vol. 4, pp. 140–149.

Thomson, W. (1990). ‘The Consistency Principle’, in (T. Ichiishi, A. Neyman and Y. Tauman, eds.), *Game Theory and Applications*, pp. 187–215, Economic Theory, Econometrics, and Mathematical Economics, San Diego: Academic Press.

Thomson, W. (2019). *How to divide when there isn't enough: from Aristotle, the Talmud, and Maimonides to the axiomatics of resource allocation*, Cambridge University Press.

UNDP (2018). ‘The Future is Decentralised’, United Nations Development Programme.

Wunder, S. and Albán, M. (2008). ‘Decentralized payments for environmental services: The cases of Pimampiro and PROFAFOR in Ecuador’, *Ecological Economics*, vol. 65(4), pp. 685–698.