Dynamic Mechanism Design for a Global Commons

Rodrigo Harrison* **and** **Roger Lagunoff†**

Final Version: March 8, 2016‡

Abstract

We model dynamic mechanisms for a global commons. Countries benefit from both consumption and aggregate conservation of an open access resource. A country's relative value of consumption-to-conservation is privately observed and evolves stochastically. An optimal quota maximizes world welfare subject to being implementable by Perfect Bayesian equilibria. With complete information, the optimal quota is first best; it allocates more of the resource each period to countries with high consumption value. Under incomplete information, the optimal quota is fully compressed — initially identical countries always receive the same quota even as environmental costs and resource needs differ later on. This is true even when private information is negligible.

JEL Codes: C73, D82, F53, Q54, Q58

Key Words and Phrases: Optimal quota, full compression, fish wars, Perfect Bayesian equilibria, international agency, climate change.

*Instituto de Economía Pontificia Universidad Católica de Chile, Av. Vicuña Mackenna 4860, Macúl, Santiago 780436. CHILE

†Department of Economics, Georgetown University, Washington, DC 20057, USA. +1-202-687-1‘510, lagunofr@georgetown.edu, www.georgetown.edu/lagunoff/lagunoff.htm.

‡We thank seminar participants at Brown, NYU, CETC Montreal, the NBER Summer Political Economy & Public Finance Workshop, USACH (Chile), University Carlo Alberto, and Masaki Aoyagi, Sandeep Baliga, Andy Postlewaite, Debraj Ray, and three anonymous referees for valuable comments and suggestions.

1 Introduction

The global effects of environmental problems such as green house gas (GHG) emissions, deforestation, and species extinction present a significant challenge for policy makers. The global scale of GHG emissions, for instance, requires that most countries be involved in any climate negotiations. Moreover, the accumulation of atmospheric GHG is an inherently dynamic process. Its effects are difficult to predict and are heterogeneous across countries. According to the Intergovernmental Panel on Climate Change (IPCC):

> "Peer-reviewed estimates of the social cost of carbon in 2005 average US$12 per tonne of CO2, but the range from 100 estimates is large (-$3 to $95/tCO2). This is due in large part to differences in assumptions regarding climate sensitivity, response lags, the treatment of risk and equity, economic and non-economic impacts, the inclusion of potentially catastrophic losses and discount rates. Aggregate estimates of costs mask significant differences in impacts across sectors, regions and populations and very likely underestimate damage costs because they cannot include many non-quantifiable impacts..." (IPCC 2007 Synthesis Report).

In light of this, any international agreement must be structured so that countries find it in their self interest to follow its prescriptions at each point in time, all the while accounting for difficult-to-predict and asymmetrically observed changes in the benefits and costs of carbon usage.

This paper examines the nature of optimal self-enforcing agreements to regulate a global commons. We posit an infinite horizon model of global resource consumption. The resource is depletable, and its aggregate use imposes environmental costs on each country. Access to the resource is not limited, and each country derives simultaneous benefit both from its own resource consumption and from the aggregate conservation of the resource stock.

Conservation is intrinsically beneficial to each country because it allows the country to avoid the environmental costs of global resource consumption. The conservation benefits are assumed to be heterogeneous across countries and are stochastically determined as countries are hit with private, idiosyncratic "payoff" shocks. The shock process in our model captures a common feature in many commons problems: environmental costs are difficult to forecast and often vary widely across countries.[1]

Our model adopts a parametric approach which, for a variety of technical reasons, has proved useful in studies in dynamic resource allocation.[2] One of the first of these was the classic common pool resource model of Levhari and Mirman (LM) (1980). LM posit a parameterized "fish war" model of an open access resource problem. In their model identical users choose how much to extract each period. The residual is left for future extraction. There are no extraction costs or associated externalities. Conservation is therefore valued in the LM model only for instrumental reasons: preserving the stock allows one to smooth consumption. The LM model and its many successors admit closed form solutions yielding a transparent view of the "tragedy of the commons" problem.

A new generation of parametric models augment the LM model. Dutta and Radner (2006, 2009) examine energy consumption with emissions externalities. Antoniadou, et al. (2013) examine resource extraction in a more general class of parametric models.

Our paper works from these blueprints, adding some modifications along the way. We consider a heterogeneous externality in resource consumption that makes conservation directly beneficial. We add private shocks that alter each country's relative value of conservation. Despite the generalization, the parameterization is tractable enough to yield explicit, close form solutions. From these explicit solutions, we examine how optimal commons mechanisms respond to uncertainty, private information, and cross-sectional variation among countries.

In particular, we solve for an *optimal quota system*. An optima quota system is an international agreement specifying each country's resource consumption (or emissions) at each date, given any carbon stock and payoff realization, such that (a) the agreement jointly maximizes the expected long run payoffs of all countries, and (b) the agreement can be implemented by a Perfect Bayesian equilibrium (PBE).

[1]Burke, et. al. (2011) find, for example, widely varying estimates of the effect of climate change on US agriculture when climate model uncertainty is taken into account. Desmet and Rossi-Hansberg (2013) quantify cross country variation in a calibrated model of spatial differences in welfare losses across countries due to global warming. These differences come primarily from geography but may be amplified by trade frictions, migration, and energy policy.

[2]See Long (2011) for a survey of the vast literature since the 1980.

Implementability in PBE captures the fact that there is no explicit mechanism designer in global commons problems. There is no global government that can impose its will on the countries. Instead, the PBE concept requires the quota system to be *dynamically self-enforcing* in compliance and in public disclosure of information.

In the benchmark case of full information, each country's payoff for consumption and conservation is known and there are no shocks. In this case, the first-best quota is shown to be implementable and is characterized by stationary extraction rates that vary across countries. Those countries that place high value on consumption (or low value on conservation) are permitted to extract more. Because the effects of full depletion are catastrophic, cheating is deterred by graduated punishments that further deplete the stock each time a country violates its prescribed resource use. For this reason, implementability does not depend on discounting.[3]

A second set of results pertain to the case of *incomplete* information — the case where persistent, private payoffs shocks hit each of the countries. Under private information, all countries have incentives to choose extraction policies that overstate their values for extraction. Hence, the first-best quota system is not incentive compatible.

Unlike the full information benchmark, we show that with private shocks the optimal quota is completely insensitive to a country's realized type. We refer to this as the property of *full quota compression*.

To illustrate what full compression means, consider two ex ante identical countries. Suppose the realized shocks are such that one country ends up with high resource needs and/or low environmental damage, while the other ends up with low resource needs and/or high environmental damage. Full compression then implies that the same quota is assigned to each country at every point in time, regardless of the initial realization of the shocks.

Full compression also holds for arbitrarily small amounts of private information. Specifically, as long as the support remains fixed, the distribution can place arbitrarily large mass on a single resource type. The result suggests that arbitrarily small amounts of private information can have first order implications for international agreements.

[3]The punishment scheme can, in principle, implement any feasible payoff.

The basic intuition for the compression results is straightforward. Global commons problems entail free riding. Under the optimal quota, all countries have individual incentives to over-extract. Free riding incentives are, of course, higher for higher types. Nevertheless, there is a threshold extraction rate such that optimal extraction rates for *all* types will lie below this cutoff. Meanwhile, free riding incentives, if left unchecked, will push countries to extract above the cutoff level. Hence, a quota that is not compressed will allow some types to indirectly free ride by mimicking other types that are allocated higher extraction quotas.

It is worth noting that this intuition applies to commons problems but not necessarily to markets. Optimal dynamic mechanisms for firms, for instance, will typically not be compressed because allocation of market share inherently requires hard trade-offs between market shares of different types of firms. A higher production quota assigned to one type of seller must be offset with a lower one to another.[4]

By contrast, in the global commons problem there are no such trade offs. A country's "market share" is its expected net present value of "stored resource" which behaves like a public good. Thus, a planner increases each country's value of stored carbon simultaneously by reducing everyone's quota. This is not a good thing for the planner, however, since a reduction of the quota increases all countries' incentives to free ride via manipulation of information. The planner is hamstrung by having no additional instrument beyond the quota to dampen these incentives.

For more clarity on this point, we consider an extension of the model in which utility of the representative agents of each country is transferable. The presence of transferable utility (TU) is shown to eliminate the inefficiency because the transfers comprise a common unit of account from which high type countries can be compensated by low types for truthful disclosure. This additional "degree of freedom" admits an equilibrium that implements the first best quota.

We argue, however, that transfers in a non-transferable utility (NTU) framework is a more natural benchmark in our setting. With NTU transfers a welfare improving scheme may exist in the context of the model, but requires countries that end up with

[4]There are exceptions. For instance Athey and Bagwell (2008) show that in certain types of markets with persistent private shocks, the optimal collusive mechanism assigns a market share to each firm that is independent of the firm's realized cost type. We discuss differences and similarities between market mechanisms and the commons mechanisms in the upcoming Literature Section and in Section 5.3

low usage value be subsidized by those with high usage value. This will sometimes require that developing countries subsidize developed ones.

Up next, Section 2 summarizes the literature on dynamic mechanisms design as it applies to global commons. Section 3 describes the benchmark model of full information. In that model there are no shocks and each country's resource type is common knowledge. Section 4 introduces private persistent shocks. Section 5 explains the logic and implications of the full compression result. We examine differences in implications if there model were to admit transferable utility, imperfectly persistent shocks, or a more general parametrization. Section 6 concludes with a discussion of potential implications for policy. Section 7 comprises an Appendix with the proofs.

2 Related Literature

There is, by now, a large literature analyzing mechanisms to address global commons problems. Understandably much of the literature focusses on a fairly narrow range of practical options. These include variations of cap and trade, carbon taxes, credit exchanges, and other well publicized proposals.[5]

A large quantitative literature has emerged to evaluate these. Some key quantitative assessments of carbon tax policies, for instance, include Nordhaus (2006, 2007) and Stern (2006), Golosov et al (2014), and Acemoglu et al (2012). Krusell and Smith (2009) calibrate a model of the global economy with fossil fuel use. They provide quantitative assessments of carbon taxation and cap and trade policies with the goal of achieving a zero emissions target. Rouillon (2010) proposes a competitive pricing scheme with directed transfers between individuals in a common pool problem.

Most of these papers focus on pricing mechanisms designed for firms and consumers in a competitive resource market.[6] The present study, however, is concerned primarily with incentives of strategic players. In that vein, Barrett (2003) and Finus (2001) argue that any international mechanism must be dynamically self-enforcing for large stakeholders. Consequently, they propose repeated game models in which

[5]See Arava, et. al. (2010) for a summary of the mechanisms in place and their rationales.
[6]See Bodansky (2004) for a summary of the hurdles faced by these proposals.

international climate agreements are implementable in subgame perfect equilibria.

A number of models in the literature extend the self-enforcement constraint to non stationary commons games that better characterize the dynamics of resource use.[7] Cave (1987) examines the traditional full information common pool model of Levhari and Mirman (1980). He shows that punishment strategies that trigger the Markov Perfect equilibrium (the "business-and-usual" outcome) can enforce full cooperation of the agreement when the participants are sufficiently patient.

Dutta and Radner (2004, 2006, 2009) and Antoniadou, et al (2013) study innovations to the LM framework. Dutta and Radner characterize the optimal emissions quota and, like Cave (1987), use the Markov Perfect benchmark as a type of Nash reversion strategy to sustain the optimum. Antoniadou, et al. focus attention on Markov Perfect equilibria in a more general, yet tractable, parametric model.

Battaglini and Harstad (2012) examine endogenous coalition formation in environmental agreements. In their model, global pollution can be addressed by investment in green technologies. The problem is that if agreement is "contractually complete" then countries may refuse to participate. Whereas if the agreement is incomplete, then it gives rise to an international hold up problem which, fortunately, can be mitigated when large coalitions of countries sign on to the appropriately structured agreement.

Like these, we stress the dynamic self-enforcement requirement of any international agreement. The inclusion private valuations for the global commons differentiates our model from the models highlighted above. With private information, dynamic self-enforcement must apply both to compliance and to truthful disclosure in any regulatory mechanism.[8]

In non-commons environments, dynamic models with private shocks are more common, and our study builds on them to some extent. For instance, Athey, Bagwell, and Sanchirico (2001), Aoyagi (2003), and Skrzypacz and Hopenhayn (2004) model optimal collusion of firms with private, iid shocks. Athey and Bagwell (2008) who

[7]Ostrom (2002) provides a broad but informal discussion of the problems involved in extending her well known "design principles" set forth in Ostrom (1990) to the global commons. See also Haurie (2008) for a survey that includes cooperative game theoretic concepts.

[8]See, for instance, Baliga and Maskin (2003) who study a static model of environmental externalities with private information.

study a model with persistent shocks. Compression-like results hold in these models when shocks are perfectly persistent and hazard rates are monotone.[9] The fact that compression holds in our paper without the monotone hazard assumption reflects a basic difference between the commons problem and collusion/auction environments. We discuss these differences in Section 5.3.

Other models with persistent shocks include Pavan, A., I. Segal, and J. Toikka (2014), and Halac and Yared (2014). The latter demonstrate optimal rules that are explicitly uncompressed and, in fact, history dependent. Their setting - a model of dynamically inconsistent government - is quite different from either collusion or commons environments.

3 The Full Information Benchmark Model

3.1 Basic Setup

This section sets up full information model as a benchmark. The model consists of n countries, indexed by $i = 1, \ldots, n$, in an infinite horizon, $t = 0, 1, \ldots$. Each country's economy makes essential use of an open access resource each period. Countries make inter-temporal strategic decisions regarding how much of the resource to extract and use.

To better motivate the framework, we use carbon usage as a leading, albeit imperfect, example.[10] The current stock of the resource at date t is given by ω_t. In the case of carbon, the current stock ω_t is the amount of "stored" greenhouse gas — the amount of carbon currently preserved under ground or in forest cover. Initially, we assume that the stock is known, and each country is able to precisely control its

[9] For related results, see McAfee and McMillan (1983) who model static collusion environments, Amador, Angeletos, and Werning (2006) who model consumption-saving environments with private, iid shocks, and Athey, Atkeson, and Kehoe (2005) who model optimal discretion by a monetary authority who privately observes iid shocks.

[10] Strictly speaking, fossil fuel is a leading source of GHG emissions and would not be characterized as a pure, open access resource. Nevertheless, we maintain the open access assumption in our discussion of carbon mechanisms because access to all types of resources that produce GHG emissions are widely dispersed among a large collection of countries. The open access model focuses attention on many of the critical difficulties in controlling GHG emissions, namely, free riding incentives, heterogeneity, and potential misrepresentation of information.

internal resource usage. We interpret the stock as a "sustainability" bound rather than an absolute quantity available. Fix the initial stock at $\omega_0 > 1$.

Country i's resource consumption at date t is c_{it}. Total consumption across all countries is $C_t = \sum_i c_{it}$. Feasibility requires $C_t \leq \omega_t$. We assume that resource use and emissions are linearly related so that $\omega_t - C_t$ of the resource remains as, for instance, the amount of stored carbon at the end of the period. The resource extraction technology is given by

$$\omega_{t+1} = (\omega_t - C_t)^\gamma \tag{1}$$

When $\gamma \leq 1$ the post-extraction stock depreciates exponentially at rate γ. However, $\gamma > 1$ allows for growth in the stock.

Let $\mathbf{c}_t = (c_{1t} \ldots, c_{nt})$ denote the date t profile of resource consumption. A country's flow payoff in date t is given by

$$\theta_i \log c_{it} + (1 - \theta_i) \log(\omega_t - C_t). \tag{2}$$

The value θ_i is the weight given to country i's log consumption, whereas $1 - \theta_i$ is the weight assigned to the remaining resource stock $\omega_t - C_t$. The parameter θ_i is country i's "resource type" or simply its "type" and is assumed to lie in an interval $[\underline{\theta}, \overline{\theta}] \subset [0, 1]$. A *type profile* is given by $\theta = (\theta_1, \ldots, \theta_n)$. Following convention, $\theta_{-i} = (\theta_j)_{j \neq i}$.

Flow payoffs are discounted by δ each period. The entire dynamic path profile of resource consumption is the given by $\mathbf{c} = \{\mathbf{c}_t\}_{t=0}^{\infty}$. A path $\mathbf{c}$ is *feasible* if it is consistent with the technological constraint (1) and $C_t \leq \omega_t$ at each date t. Given a feasible consumption path $\mathbf{c}$, the long run payoff to country i at date t may be expressed recursively as

$$U_i(\omega_t, \mathbf{c}, \theta) \;=\; \theta_i \log c_{it} + (1 - \theta_i) \log(\omega_t - C_t) \;+\; \delta\, U_i(\omega_{t+1}, \mathbf{c}, \theta). \tag{3}$$

3.2 Interpretation and Technical Issues in Commons Models

Using carbon as an example, a country's payoff U_i in (3) can be interpreted in one of two ways. The first way is to associate U_i simply with the preferences of a "representative citizen." The citizen's flow payoff weights both resource consumption and resource conservation. Since the costs of GHG emissions are associated with consumption of carbon-based resources, the citizen therefore derives some value from keeping the carbon in its "stored" state.

The "pure preference interpretation" builds on, and may be compared to, traditional "fish war" models of common pool resource usage dating back to Levhari and Mirman (1980). Those models assume $\theta_i = 1$ for all i, in which case a user of the resource merely trades off the value of present with future usage, given the anticipated usage of others. A user's value of " conservation" in the traditional model is therefore purely instrumental. Conservation is valued because it represents potential future usage, and the user prefers to smooth consumption.

The present formulation differs by adding a direct preference for resource conservation. This preference, moreover, is heterogeneous across countries.[11] In the case of carbon-based resources, countries obviously value the use of fossil fuels and timber, but recognize the associated GHG emissions as a costly by-product. Both benefits and costs of usage differ across countries. Warmer average temperatures resulting from GHG emissions are viewed differently in Greenland than in Sub-saharan Africa.

A second interpretation is that θ_i reflects production elasticity of a carbon based resource. According to this "production-based" interpretation, all representative consumers have identical payoffs of the form

$$\sum_t \delta^t \log y_{it}$$

where y_{it} is a composite output consumed by the representative consumer from country i at date t, and δ is a common discount factor. The composite good is produced using two inputs, extracted and unextracted carbon, according to the technology $y_{it} = c_{it}^{\theta_i}(\omega_t - C_t)^{1-\theta_i}$.

In this formulation, the carbon extracted from the ecosystem gets used up in the production process. The unextracted or "stored" carbon from ecosystem is renewable, but depreciates/appreciates at rate γ according to (1). Each country utilizes the inputs at different intensities. Countries with larger θ_i use more of the carbon input to produce a given output. Richer countries, for instance, have larger carbon

[11]The recent models of Dutta and Radner (2006, 2009) are among the few others we are aware of that build in heterogeneous usage externalities in the common pool framework.

requirements as a consequence of a more developed economy.[12]

In either interpretation, we demonstrate in the next section how the parametric specification yields an explicit closed form solution for an optimal self-enforcing contract. This is significant since abstract dynamic game models are notoriously difficult to solve and even more difficult to interpret. Equilibrium existence itself is not guaranteed except under fairly special circumstances.[13] Even in such circumstances, it is difficult to say too much about the commons in the abstract. It is has proved difficult for instance to identify conditions on preferences and laws of motion of resource stocks that ensure well behaved value functions.[14] Without parametric structure or numerical approximations, results are sparse.[15]

For these reasons, much of the prior literature tended toward parametric approaches. The multiple interpretations above suggest that our parameterization is flexible and well suited to address the role of heterogeneity and private information inherent in global commons problems.

3.3 Optimal Quota Systems

Our interest is in international agreements chosen by the participants through a coordinating body such as the United Nations (U.N.). We refer to this body as the International Agency (IA). The IA, as envisioned here, operates by the consent of its members, gathers and makes available information, makes recommendations, and suggests sanctions for violations. It does not have the power to restrict communication or enforce sanctions.

Initially, we consider the case in which all countries' types are common knowledge and fixed throughout. Later, we consider the case of privately observed, stochastically determined types for each country i. The model "abstracts away" issues of endogenous technical change and technology transfer between countries. Though these are clearly

[12]A country's type θ_i does not necessarily correspond to its size. While larger countries would have greater need for resources, the costs of climate change may be larger as well. The country's type θ_i only determines its *relative* weight between use and conservation. Size differentials could be captured instead by differential welfare weighting in any planner's problem.

[13]See Dutta and Sundaram (1998) for a lucid discussion of the technical difficulties in establishing existence of equilibria in stochastic games.

[14]The technical problems examined in Mirman (1979) remain highly relevant.

[15]See, for instance Doepke and Townsend (2006) or Cai, Judd, and Lontzek (2012).

central issues in current discussions of climate mechanisms, the present study focuses at this stage purely on issues of disclosure and compliance.

Hence, because the IA cannot directly impose or enforce anything, it merely "recommends" a *quota system*. A quota system is defined as a mapping $\mathbf{c}^*(\theta)$ from type profiles to feasible consumption paths. Here $c^*_{it}(\theta)$ is the targeted resource consumption recommended for country i at date t given the global type profile θ. The recommended quota system must be *implementable* by a subgame perfect equilibrium (SPE) of the dynamic resource game.

In order to define the implementability requirement, a few extra bits of notation are needed. Let $h^t = (\omega_0, \mathbf{c}_0, \omega_1, \mathbf{c}_1, \ldots, \omega_{t-1}, \mathbf{c}_{t-1}, \omega_t)$ be the date t history of consumption and resource stocks, including the current stock ω_t. The initial history is $h^0 = \omega_0$. A *usage strategy* $\sigma_i(h^t, \theta)$ for country i maps histories and global type profiles to desired resource consumption c_{it} at date t. A *usage profile* is given by $\sigma = (\sigma_1, \ldots, \sigma_n)$. A default rule describing payoffs when $\sum_i c_{it} = \omega_t$ (full depletion) is needed to complete the specification of the game.[16]

A profile σ is a *subgame perfect equilibirum (SPE)* if each σ_i maximizes country i's long run payoff

(4)
$$V_i(h^t, \sigma, \theta) \;=\; \theta_i \log\left(\sigma_i(h^t,\;\theta)\right) \;+\; (1-\theta_i)\log\left(\omega_t - \sum_{j=1}^{n} \sigma_j(h^t,\;\theta)\right) \;+\; \delta\, V_i(h^{t+1}, \sigma, \theta).$$

after every history h^t. Note that in the recursive formulation in (4), decisions in t determine the history h^{t+1} entering $t + 1$.[17]

An SPE profile σ may then be said to *implement* a quota system $\mathbf{c}^*$ if σ generates consumption $\mathbf{c}^*$ along the equilibrium path. Formally, $\mathbf{c}^*$ is generated recursively: $\mathbf{c}^*_0(\theta) = \sigma(h^0, \theta)$, then $\mathbf{c}^*_1(\theta) = \sigma(h^0, \mathbf{c}^*_0(\theta), (\omega_0 - C^*_0(\theta))^\gamma,\; \theta)$, and so forth... Hence, if a SPE profile σ implements a quota system $\mathbf{c}^*(\theta)$, it follows that $V_i(h^0, \sigma, \theta) = U_i(\omega_0, \mathbf{c}^*(\theta), \theta_i)$ where V_i is defined by (4) and U_i is defined by (3).

The IA's problem can now be formally stated as one that recommends both a

[16]The rule is needed to map all strategy profiles to well defined payoffs in the game. Hence, extend the definition of the flow payoff u_i to the extended reals: let $u_i = \theta_{it}\log c_{it} + (1-\theta_{it})\log(\omega_t - C_t)$ as in Equation (2) when $c_{it} > 0$ and $C_t < \omega_t$, and let $u_i = -\infty$, otherwise.

[17]Formally, $h^{t+1} \;=\; (h^t, c_t, \omega_{t+1}) \;=\; \left(h^t,\; \sigma(h^t,\;\theta),\; (\omega_t - \sum_{j=1}^{n}\sigma_j(h^t,\;\theta))^\gamma\right)$.

quota $\mathbf{c}^*(\theta)$ and a subgame perfect profile σ^* that solve

$$\max_{\mathbf{c}^*(\theta)} \sum_{i=1}^{n} U_i(\omega_0, \mathbf{c}^*(\theta), \theta) \tag{5}$$

$$\text{such that } \mathbf{c}^*(\theta) \text{ is implemented by a } \sigma^*.$$

Hence, the IA chooses the quota to maximize the joint sum of all countries' payoffs such that the quota is sequentially self-enforcing.

The formulation in (5) implicitly assumes that all countries are of the same size. To account for size differences, an international agency would attach differential welfare weights. Doing so would complicate the notation without adding to the results.

The solution to (5) can be found by breaking the problem into two steps.

- Step 1 characterizes the optimal quota without the compliance constraints. In the full information case, this amounts to solving (5) without the SPE equilibrium constraint. This "relaxed" problem yields the unconstrained, "first-best" quota chosen by an international agency with the ability to impose and enforce its choice upon the participants.
- Step 2 shows that the first-best quota can be implemented by SPE profile σ^*.

This simple two-step algorithm will be repeated later on when private shocks are introduced into the model.

3.4 Finding the Optimal Quota

Step 1 in the above algorithm finds the optimal quota as a solution to a "relaxed problem." A candidate solution $\mathbf{c}^*$ to this "relaxed" problem (outlined in Step 1) satisfies the Bellman equation

(6)

$$\sum_i U_i(\omega_t, \mathbf{c}^*, \theta) \;=\; \max_{c_t} \left[\sum_i \theta_i \log c_{it} + (1-\theta_i)\log(\omega_t - C_t) \;+\; \delta \sum_i U_i(\omega_{t+1}, \mathbf{c}^*, \theta) \right]$$

subject to (1).

To find the solution to this Bellman equation, it's easier to work with extraction *rates* rather than levels. For any dynamic consumption path $\mathbf{c}$, let $\mathbf{e} = (e_{it})$ denote the corresponding path of extraction rates where $c_{it} = e_{it}\omega_t$. Let $\mathcal{E}_t = \sum_i e_{it}$ denote the aggregate extraction rate so that $C_t = \mathcal{E}_t\omega_t$.

Using rates rather than levels in the Bellman equation, the first order condition in e_{it} is:

$$\frac{\theta_i}{e_{it}} - \frac{\sum_j (1-\theta_j)}{1-\mathcal{E}_t} - \delta\gamma \sum_j \frac{\partial U_j(\omega_{t+1}, \mathbf{e}, \theta)}{\partial \omega_{t+1}} \, \omega_t^{\gamma}(1-\mathcal{E}_t)^{\gamma-1} = 0.$$

Appendix 7.1 has the full derivation. From the first order condition, the IA's Euler equation is

$$(7) \qquad \frac{\theta_i(1-\mathcal{E}_t)}{e_{it}} - \sum_j (1-\theta_j) = n\delta\gamma + \delta\gamma\left(\frac{\theta_i(1-\mathcal{E}_{t+1})}{e_{i\,t+1}} - \sum_j (1-\theta_j)\right).$$

This Euler equation is derived using standard techniques and the Envelope Theorem. The forward solution to the IA's Euler equation is then easily found to be

$$(8) \qquad \frac{\theta_i(1-\mathcal{E}_t)}{e_{it}} - \sum_j (1-\theta_j) = \frac{n\delta\gamma}{1-\delta\gamma}$$

which is stationary.[18] Re-arranging terms and aggregating over i yields the aggregate rate $\mathcal{E}^*(\theta) = \frac{\Theta(1-\delta\gamma)}{n}$ where $\Theta \equiv \sum_j \theta_j$. In other words, the first-best aggregate extraction rate is a fraction $(1-\delta\gamma)$ of the average resource type $\frac{\Theta}{n}$. This aggregate rate is achieved by country-specific extraction rates.

$$(9) \qquad e_i^*(\theta) = \frac{\theta_i(1-\delta\gamma)}{n},$$

which is displayed in Figure 1. By recursive substitution, the extraction rates yields an explicit closed form solution for the optimal quota system:

$$(10) \qquad c_{it}^*(\theta) = \frac{\theta_i(1-\delta\gamma)}{n} \, \omega_0^{\gamma^t}\left(1 - \frac{\Theta(1-\delta\gamma)}{n}\right)^{\gamma(1-\gamma^t)/(1-\gamma)}$$

for country i in date t.[19]

[18] I.e., it is independent of the current stock ω and of calendar time. Note that to ensure the solution exists, one must assume $\delta\gamma < 1$

[19] To obtain (10), start with $t = 0$: $c_{i0}^*(\theta) = e_i^*\omega_0 = \frac{\theta_i(1-\delta\gamma)}{n}\omega_0$, then for $t = 1$, $c_{i1}^*(\theta) = e_i^*\omega_1 = e_i^*(\omega_0 - C_0^*)^{\gamma} = \frac{\theta_i(1-\delta\gamma)}{n}\omega_0^{\gamma}(1 - \frac{\Theta(1-\delta\gamma)}{n})^{\gamma}$, ... and so on.

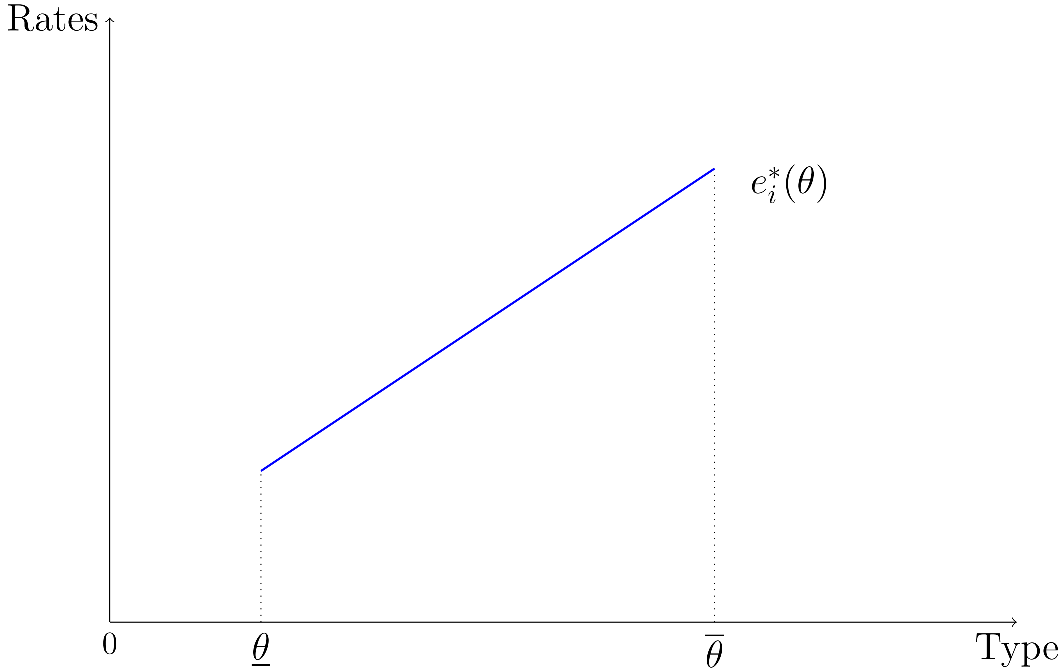

Figure 1: First Best Extraction Rates

Notice that the quota allocated to each country declines or increases over time, depending on whether the stock is exhaustible ($\gamma \leq 1$) or renewable ($\gamma > 1$). Notice also that both the rate and the level in (10) are increasing in one's own resource type θ_i but decreasing in the cross-country average Θ/n. Hence, countries with larger-than-average carbon usage intensity can extract more. Put another way, "pro-consumption" types should extract more while "pro-conservation" types should extract less.

3.5 Implementing the Optimal Quota

As for Step 2, compliance with the first best quota is not automatic. One could compare the IA's Euler equation to that of each individual country. That latter corresponds to a country's individual incentives for resource consumption, and is given by:[20]

$$\frac{\theta_i(1-\mathcal{E}_t)}{e_{it}} - (1-\theta_i) \;=\; \delta\gamma + \delta\gamma\left(\frac{\theta_i(1-\mathcal{E}_{t+1})}{e_{i\,t+1}} - (1-\theta_i)\right).$$

[20]The derivation mirrors the algorithm for solving the planner's problem in Appendix 7.1.

As before, the forward solution is easily found.

$$\text{(11)} \qquad \frac{(1-\mathcal{E}_t)}{e_{it}} = \frac{\delta\gamma}{\theta_i(1-\delta\gamma)} + \frac{1-\theta_i}{\theta_i}$$

From (11), one can calculate country's best response to the extraction rates $\mathcal{E}_{-it}$ of other countries:

$$\text{(12)} \qquad e_{it} \;=\; BR_i(\theta_i, \mathcal{E}_{-it}) \;=\; \theta_i(1-\delta\gamma)(1-\mathcal{E}_{-it})$$

Predictably, pro-extraction types have a greater incentive to extract more. It is easy to verify that $BR_i(\theta_i, \mathcal{E}^*_{-i\ t}) > e^*_i$. In other words, the individual country's incentives are toward greater extraction than that prescribed by the IA.[21]

Consequently, the quota system described by (10) is the optimal one only if it can be implemented by a SPE (Step 2 in the solution algorithm). This is summarized in the following result.

Proposition 1 (Full Information Benchmark) *Consider the model with full information satisfying $\delta\gamma < 1$. Then the IA's optimal quota system (the solution to (5)) is the quota system $\mathbf{c}^*(\theta)$ described by (10). Namely,*

$$\text{(13)} \qquad c^*_{it}(\theta) \;=\; \frac{\theta_i(1-\delta\gamma)}{n}\,\omega_0^{\gamma^t}\left(1-\frac{\Theta(1-\delta\gamma)}{n}\right)^{\gamma(1-\gamma^t)/(1-\gamma)} \qquad \forall\, i \;\; \forall\, \theta \quad \forall\, t$$

Proposition 1 is a useful benchmark for comparing subsequent results with private information. Note that the implementation result holds for *any* discount factor $\delta > 0$, including impatient ones. This is possible because payoffs are unbounded below, reflecting the idea that in a global commons, the costs of full resource depletion may be catastrophic.

As a consequence, the IA can recommend further threats of resource depletion in any continuation payoff, even ones that are already punitive, to enforce compliance. Since increased resource depletion hurts all countries, the credibility of the punishment

[21] Note that solving (12) as a simultaneous system of equations across countries yields the unique Markov Perfect equilibrium (MPE) — the analogue of the "business-as-usual equilibrium" computed by Cave and by Dutta and Radner in their models. The analytical solution to the MPE in the present model is worked out in Harrison and Lagunoff (2015, Sect. 4).

depends on even harsher punishment if the countries fail to carry out the sanction. This means, in turn, that increasingly severe depletion threats must be used, each such threat made credible only by even more severe depletion threats later on, and so forth. The successive threats are only reached, of course, by further deviations at each counterfactual stage. Since each country's payoff in the residual stock is unbounded, the sequence of threats can be recursively defined. The proof in the Appendix gives the formal details.

Further, the equilibrium does not actually require monitoring by the IA, since the punishments at each counterfactual stage are not tailored to the perpetrator who deviated from the prescribed rate. Instead, it need only monitor the aggregate stock itself to determine whether a deviation occurred.

Finally, we emphasize that the algorithm for constructing punishments is not sensitive to the location of the social optimum. Indeed, because the punishments are potentially unbounded, any finite payoff profile can be implemented in this way.

3.6 Credibility of the Implementation

Given the construction, a few questions naturally arise. First, does implementation works in a commons model for any flow payoff that diverges to $-\infty$ as $\mathcal{E}_t \to 1$? The proof in the Appendix indicates the answer is yes, although we do not offer an explicit proof of the general case.

Second, is it credible for the IA to carry out draconian punishments in the event of repeated violation? One can argue this both ways. We point out that the IA here is merely a coordination device — it has no enforcement powers. Thus the punishments are self-enforcing, carried out as they are by the participants, and sustained by the self-fulfilling expectations of others' behavior. In this sense, the international agency need not be credible in order for the implementation to work.

Third, how would implementation proceed if unbounded punishments were not permissible or feasible? The most direct alternative to the present construction is a Markov Perfect reversion strategy. The idea is straightforward. If a country defects from the optimal quota, then all countries revert to a Markov Perfect Equilibrium from that period forward. The deterrent works if the participants are patient enough.

This type of implementation is described in Cave (1987), Dutta and Radner (2006), and in a more closely related model of ours (Harrison and Lagunoff (2015), Section 4). Significantly, even when participants are patient, the reversion strategy only works if the participants are not too heterogeneous.[22]

4 Persistent Private Shocks

This Section investigates the nature of optimal quotas when each country's internal costs and benefits of resource usage are privately observed as countries are hit with idiosyncratic private shocks.

The shocks capture a degree of unpredictability of the effect of climatic change within each country. Moreover, country's resource shock is assumed to be observable only to that country, presuming that countries have inside knowledge of changing business conditions and "local inventories" of sources and sinks of GHG.

Consider a scenario, for instance, in which a warmer climate lowers a country's agricultural yields. This, in turn, leads to more intense use of petroleum-based fertilizer, thus increasing the country's relative value of carbon. In another scenario, the warmer climate leads to an increase in the saline contamination of the country's fresh water fisheries. In that case, the relative value of carbon decreases as GHG costs increase. The international agency is keen to know which scenario prevails, but must rely on self-reported information by the countries, a fact explicitly recognized in many international agreements. See, for instance, Article 12 under the UN Framework Convention for Climate Change, and Articles 5 and 7 under the Kyoto Protocol.

The shock is assumed to hit i's type at $t = 0$. One natural interpretation is that this is the limiting case of a more general scenario in which each country incurs serially correlated private shocks to its resource type each period. The present model corresponds to the case in which shocks are perfectly persistent.[23] Perfect persistence is a reasonable approximation to a situation where environmental change is more

[22] The problem is that if a country's type θ_i is too low, then it may prefer the MPE over a planner's solution with equal welfare weights.

[23] An external appendix explore the general case and extends our main result (Proposition 2) to this case. See faculty.georgetown.edu/lagunofr/optimal-resource10-External-Appendix2.pdf.

rapid than technological progress.

Country i's type θ_i is therefore determined by a distribution $F_i(\theta_i)$. For simplicity we analyze a symmetric situation in which $F_i = F_j$ for all i and j. Under the production interpretation of the model the shocks to types represent exogenous technological change. We assume the shocks are independent across countries, with $F_i(\cdot)$ differentiable in θ_i, has full support on $[\underline{\theta}, \overline{\theta}]$, and admits a continuous density $f_i(\cdot)$. The distribution F_i is assumed to be commonly known to all countries, however, each country's realized shock each period is privately observed.

4.1 Optimal Quota Systems with Private Shocks

As before, the international agency (IA) recommends a quota system. To make an effective recommendation, the international agency is a gatherer and dispenser of information. The IA solicits information concerning each country's realized type θ_i. As before, the IA is a weak mechanism designer. It cannot compel the countries to report truthfully. Nor can it withhold information since the revelation of information by each country is public. Instead, the IA serves only as a vehicle for coordinating information and usage. Again, this is appropriate in the international setting where there is no global government with external enforcement capability.

Each member country chooses whether or not to disclose its type (as, for instance, when countries make public their national income accounts, estimates, and forecasts). A country's reported type is denoted by $\tilde{\theta}_i$. The entire profile of types reported by all countries is $\tilde{\theta} = (\tilde{\theta}_1, \ldots, \tilde{\theta}_n)$.

A quota system then conditions on the reported profile $\tilde{\theta}$. To distinguish this from the full information case, we denote a quota system here by $\mathbf{c}^\circ$. Formally, a quota system is given by the sequence $\mathbf{c}^\circ = \{c_t^\circ(\tilde{\theta})\}, t = 0, 1, \ldots$.

To obtain the desired quota system, the IA recommends both a usage and a disclosure strategy. Formally, a *disclosure strategy* for country i is map $\mu_i(\theta_i)$ describing i's report at the beginning of $t = 0$ after realizing its type.

A usage strategy now is a map $\sigma_i(h^t, \tilde{\theta}, \theta_i) = c_{it}$ determining i's consumption at date t given the extraction history, the disclosure profile, and its resource type.

To an extent, the disclosure mechanism resembles some existing protocols. Article 12 of the UN Framework Convention on Climate Change, for instance, requires its signatories to periodically submit, among other things, a "national inventory of anthropogenic emissions," and a "specific estimate of the effects that the policies and measures ... will have on anthropogenic emissions by its sources and removals by its sinks of greenhouse gases..."

Let $\mu = (\mu_1, \ldots, \mu_n)$ and $\sigma = (\sigma_1, \ldots, \sigma_n)$ denote profiles of disclosure and resource usage, resp., and after any history, let $\mu(\theta_i) = (\mu_i(\theta_i))_{i=1}^n$ and $\sigma(h^t, \tilde{\theta}, \theta) = (\sigma_i(h^t, \tilde{\theta}, \theta_i))_{i=1}^n$. Given a strategy pair (σ, μ), the long run expected payoff to a country i at the resource consumption stage in date t is

(14)

$$V_i(h^t, \tilde{\theta}, \sigma, \mu|\ \theta_i) \ \equiv \int_{\theta_{-i}} \Bigg[\theta_i \log \sigma_i(h^t, \tilde{\theta}, \theta_i) + (1-\theta_i)\log(\omega_t - \sum_{j=1}^{n} \sigma_j(h^t,\ \tilde{\theta}, \theta_j)\ ) + \delta V_i(h^{t+1}, \tilde{\theta}, \sigma, \mu|\ \theta_i)\Bigg]\ dF^{\circ}_{-i}(\theta_{-i}|\ h^t, \tilde{\theta})$$

where $F^{\circ}_{-i}(\theta_{-i}|\ h^t, \tilde{\theta})$ will denote the posterior update about other countries' resource types when h^t is the usage history, $\tilde{\theta}$ is the disclosure profile, and (implicitly) given the strategy pair (σ, μ).

At the disclosure stage, country i's interim payoff before observing the disclosed type of others is

$$\int_{\theta_{-i}} V_i(h^0, \mu(\theta), \sigma, \mu|\ \theta_i) dF_{-i}(\theta_{-i}).$$

To implement a quota system, the IA recommends a profile (μ, σ) of disclosure and usage strategies. Each period it solicits information from each country about its type. If these prescriptions are followed, then all countries disclose their types according to μ. The IA then makes public the reported profile $\tilde{\theta}$. We focus on truth-telling disclosure strategies, i.e., those in which μ prescribes $\tilde{\theta}_i = \theta_i$ for each country.

The strategy pair (μ, σ) with truth-telling disclosure may then be said to *implement* the quota system $\mathbf{c}^{\circ}$ in the private shocks model if (μ, σ) yields $\mathbf{c}^{\circ}$ along the outcome path.[24]

[24] As with the full information case, the consumption path is generated recursively: $\mathbf{c}^{\circ}_0(\theta) = \sigma(h^0, \theta,\ \theta)$, then $\mathbf{c}^{\circ}_1(\theta) = \sigma(h^1, \theta)$, where $h^1 = (h^0, \sigma(h^0, \theta,\ \theta), (\omega_0 - C_0(\theta))^{\gamma}\ )$, etc.

A quota system is feasible only if it can be implemented by, a Perfect Bayesian equilibrium strategy pair (σ, μ). The pair (μ, σ) and a belief system $F_i^\circ(\theta_i | \, h^t, \tilde{\theta}), \; i = 1, \ldots n$ constitute a *Perfect Bayesian equilibrium (PBE)* if (i) at the consumption stage in date t, σ_i and μ_i together maximize i's long run expected payoff (defined in (14)) given usage history h^t, disclosure profile $\tilde{\theta}$, i's type θ_i, and given the strategies of other countries; (ii) at the disclosure stage, σ_i and μ_i together maximize i's expected payoff given h^t, given θ_i, and given the strategies of others; and (iii) beliefs F° satisfy Bayes' Rule wherever possible.

Note that at the disclosure stage, countries contemplate disclosure deviations from the prescribed plan, taking account of the fact that they have the freedom to deviate in their extraction plans at a subsequent stage. This potential for "thoughtful" deviations limits the types of punishments that any IA can suggest to the members. This also complicates the members' beliefs off-path. After any deviation from prescribed usage strategies, other countries must determine what type of deviation — a resource use deviation, or an earlier disclosure deviation, or both — occurred.

If, however, (σ, μ) is a truth-telling PBE that implements $\mathbf{c}^\circ$, then along the equilibrium path, the realized payoff for i following any type history θ^t satisfies

$$V_i(h^t, \theta, \sigma, \mu | \, \theta_i) \; = \; U_i(\omega_t, c^\circ(\theta), \theta_i) \tag{15}$$

Consequently, the IA will recommend a quota system $\mathbf{c}^\circ$ and a PBE (σ, μ) that solves

(16)

$$\max_{\mathbf{c}^\circ} \; \sum_i \; \int_\theta U_i(\omega_0, c^\circ(\theta), \theta_i) dF(\theta) \quad \text{subject to } \mathbf{c}^\circ \text{ implemented by the PBE } (\sigma, \mu).$$

Any quota that solves (16) will be called an *optimal quota*. Our main result is that the optimal quota has a special form which we refer to as *fully compressed*. A quota system $\mathbf{c}^\circ$ is fully compressed if for every country i, the recommended quota c_{it}° for country i at date t does not vary with the realized profile of shocks θ. In other words, the quota is completely insensitive to the countries' realized preferences/ production intensities for carbon.

4.2 A Full Compression Result

For the result below, one further assumption is required:

Dispersion Restriction. *The parameters* n, δ, γ, *and* $\underline{\theta}$ *jointly satisfy* $\underline{\theta} > \frac{1}{1+(n-1)\delta\gamma}$.

The dispersion restriction is a weak condition that bounds the relative dispersion of types by the discounted number of participants.[25] In general the larger the global commons problem (the larger is n), the more dispersed the resource types can be and still satisfy the constraint.

Proposition 2 *Consider the model with private, perfectly persistent shocks such that the dispersion restriction is satisfied. Then the optimal quota* $\mathbf{c}^\circ$ *(the solution to (16)) is fully compressed. Specifically,*

(1) the extraction rates corresponding to $\mathbf{c}^\circ$ *are stationary, fully compressed, and given by:*

$$e_i^\circ = \frac{(1-\delta\gamma)\left[\int_{\underline{\theta}}^{\overline{\theta}} \theta_i dF_i(\theta_i)\right]}{n} \tag{17}$$

for each country i.

(2) $\mathbf{c}^\circ$ *is implementable by an equilibrium in which each country's prescribed extraction rate after any history is stationary and fully compressed.*

To illustrate what compression means for the optimal quota, compare the optimal rate $\mathbf{e}^\circ$ in (17) to the full information solution in (9). This is illustrated in Figure 2. Under private shocks, the optimal quota assigns all countries *identical* extraction rates, despite the fact that the international agency can condition its recommendation on the information disclosed by each of the countries (again see Figure 2). If, for instance, F_i is uniform on $[\underline{\theta}, \overline{\theta}]$, then in the optimal quota, all countries are required

[25] To illustrate just how weak it is, consider $\delta\gamma = .8$ and $n = 193$ (the number of countries in the U.N.). Then the lower bound $\underline{\theta}$ need only exceed .0065 (approximately) to satisfy the dispersion constraint. To take an even more pessimistic case, consider $\delta\gamma = .5$ and $n = 34$ (the number of counties in the OECD). Then the dispersion constraint is satisfied when $\underline{\theta} > .0571$.

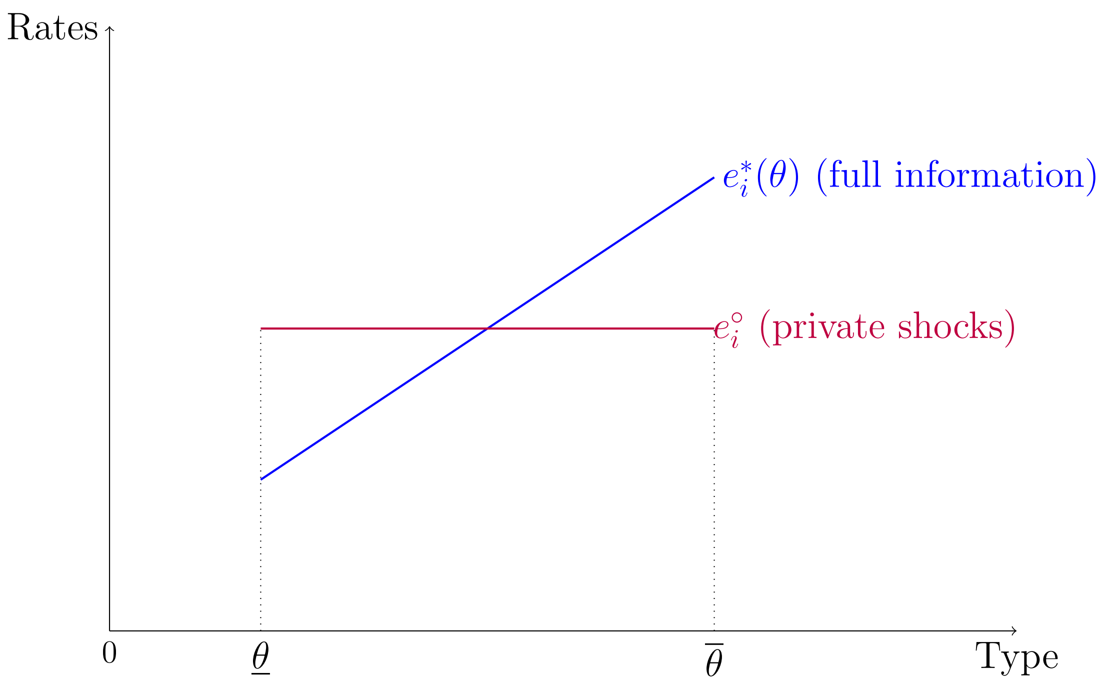

Figure 2: Optimal quota: Private Shocks vs Full Information

to extract resources at the same rate $e_i^\circ = (1-\delta\gamma)(\overline{\theta}+\underline{\theta})/2n$, even if they realize very different resource types.

Proposition 2 has troubling implications for any prospective climate agreement. Countries with realized usage values above the mean must extract less than under the full information optimum. Those below the mean can extract more (Fig. 2). Generally, the informational rents accorded to low types gives them considerable "bargaining power." Relative to the full information optimum, high types subsidize low types. In concrete terms, it suggests that fast-developing countries with higher than expected resource demand (India, Brazil, and China) must, in a sense, subsidize countries with lower than expected resource demand (U.S., Japan, EU countries).

Moreover, the compression holds even if private information on the country shocks is negligible. To see this, fix any ϵ-neighborhood of a particular type $\hat{\theta} \in [\underline{\theta}, \overline{\theta}]$ and consider a distribution F_i that places all but ϵ of its mass on that neighborhood while maintaining full support. Observe then that Proposition 2 holds for any $\epsilon > 0$.

It is worth noting that while the Proposition describes the optimal quota by its extraction rates, one can characterize the optimal consumption profile by

$$(18)\quad c_{it}^{\circ}(\theta) \;=\; \frac{(1-\delta\gamma)\int_{\underline{\theta}}^{\overline{\theta}}\theta_i dF_i(\theta_i)}{n}\;\omega_0^{\gamma^t}\left(1-\frac{(1-\delta\gamma)\sum_i\int_{\underline{\theta}}^{\overline{\theta}}\theta_i dF_i(\theta_i)}{n}\right)^{\gamma(1-\gamma^t)/(1-\gamma)}$$

for each country i and each date t. Here, consumption is not stationary but it is compressed. Consumption may grow or contract over time depending on whether the depletion rate γ admits growth.

Before proceeding with the logic of Proposition 2, we consider what the planner would choose if it *never* observed the realized types. In that case, the optimal mechanism is, of course, compressed and coincides with the compressed mechanism in the Proposition.

5 The Logic of Full Compression

5.1 Intuition from first order conditions

This section illustrates the basic logic of the result.[26] The proof proceeds by first introducing a "relaxed planner's problem" similar to the one in the full information benchmark. In this case, the PBE constraint in the relaxed problem is replaced by a weaker constraint that requires only truthful disclosure. Compliance incentives following disclosure are ignored. The International Agency then faces only a single disclosure stage at the beginning of $t = 0$ (since compliance is ignored). This is given by the truth-telling constraint:

$$(19)\qquad \int_{\theta_{-i}} U_i(\omega, c^{\circ}(\theta), \theta_i)dF_{-i} \geq \int_{\theta_{-i}} U_i(\omega, c^{\circ}(\tilde{\theta}_i, \theta_{-i}), \theta_i)dF_{-i} \quad \forall\theta_i \;\; \forall\tilde{\theta}_i$$

The "relaxed planner's problem" is stated as

$$(20)\qquad \max_{\mathbf{c}^{\circ}} \sum_i \int_{\theta} U_i(\omega, c^{\circ}(\theta), \theta_i)dF(\theta) \quad \text{subject to (19).}$$

To understand the argument, we consider, as a solution to the relaxed problem, an extraction plan $\mathbf{c}^{\circ}$ that is both stationary and symmetric. This is later established in the proof. For now, these properties are taken as given.

[26] A more detailed outline of the proof is given in Section 7.3.1.

To show how non-compression is ruled out, observe that for any profile θ_{-i} of others' types, the full information optimal extraction rate for country i possessing the highest type $\overline{\theta}$ is

$$e_i^*(\overline{\theta}, \theta_{-i}) \;=\; \overline{e} \;\equiv\; \overline{\theta}(1-\delta\gamma)/n.$$

This extraction rate serves as a cutoff. The proof shows that solution to the indirect problem must be consistent with an extraction plan $e_i^\circ(\theta)$ for each i that lies *below* $\overline{e}$. By contrast, under the Dispersion restriction, if a type could freely choose its extraction rate e_i then it would choose a larger extraction rate than $\overline{e}$. This is shown formally in the proof, but can be illustrated heuristically using first order conditions.

First, rewrite the payoffs. One can verify that under any stationary plan, country i's payoff in the relaxed problem may be expressed as

$$U_i(\omega_0, \mathbf{c}^\circ(\theta), \theta_i) \;=\; k_0 \;+\; k_1 \log(1-\mathcal{E}^\circ(\theta)) \;-\; \theta_i k_2 \log\left(\frac{1-\mathcal{E}^\circ(\theta)}{e_i^\circ(\theta)}\right) \tag{21}$$

where $\mathbf{e}^\circ$ is the stationary extraction rate corresponding to $\mathbf{c}^\circ$, and k_0, k_1, k_2 are positive constants given by

$$k_0 = \frac{\log \omega_0}{1-\delta\gamma}, \quad k_1 = \sum_{t=0}^{\infty} \delta^t \sum_{j=0}^{t} \gamma^{t-j}, \quad \text{and} \quad k_2 = \frac{1}{1-\delta} \tag{22}$$

We illustrate why separating equilibria cannot arise in the case where extraction plan e_i° for country i is differentiable in θ_i on some subinterval $\Theta_i \subseteq [\underline{\theta}, \overline{\theta}]$. Using (21), the interim payoff to type is θ_i is

$$k_0 \;+\; \int_{\theta_{-i}} \left(k_1 \log(1-\mathcal{E}^\circ(\theta)) \;-\; \theta_i k_2 \log\left(\frac{1-\mathcal{E}^\circ(\theta)}{e_i^\circ(\theta)}\right)\right) dF_{-i}(\theta_{-i})$$

Consequently, the incentive constraint for an interior type $\theta_i \in \Theta_i$ satisfies the first order condition,

$$\int_{\theta_{-i}} \left[\frac{k_2\theta_i}{e^\circ(\theta)} \;-\; \frac{k_1 - k_2\theta_i}{1-\mathcal{E}^\circ(\theta)}\right] \frac{\partial e_i^\circ(\theta)}{\partial \theta_i} dF_{-i}(\theta_{-i}) \;=\; 0$$

If e° is fully compressed, then $\frac{\partial e_i^\circ(\theta)}{\partial \theta_i} = 0$ and the first order condition is trivially satisfied. Suppose that for some type θ_i, the set $\Theta_{-i}(\theta_i) \equiv \{\theta_{-i} : \frac{\partial e_i^\circ(\theta)}{\partial \theta_i} \neq 0\}$ has F_{-i}-positive probability. Then the first order condition becomes

$$\int_{\theta_{-i} \in \Theta_{-i}} \left[\frac{k_2\theta_i}{e^\circ(\theta)} \;-\; \frac{k_1 - k_2\theta_i}{1-\mathcal{E}^\circ(\theta)}\right] \frac{\partial e_i^\circ(\theta)}{\partial \theta_i} dF_{-i}(\theta_{-i}) \;=\; 0$$

One can contrast this with first order condition for i's best response, call it e_i^B, to the relaxed solution e_{-i}° under full information:

$$\frac{k_2\theta_i}{e_i^B(\theta)} \; - \; \frac{k_1 - k_2\theta_i}{1 - \mathcal{E}_{-i}^\circ(\theta) - e_i^B(\theta)} \; = \; 0 \quad \forall \, \theta.$$

We can then use the dispersion restriction to show that any solution $\mathbf{e}^\circ$ to the relaxed problem satisfies $e_i^\circ(\theta) < e_i^B(\theta)$ for each i and almost every θ. In other words, if it were possible for a country to freely choose its extraction rate e_i, it would opt for one larger than allowed by the relaxed optimal quota, regardless of the realized θ. From this free rider problem, it follows that

$$\frac{k_2\theta_i}{e_i^\circ(\theta)} \; - \; \frac{k_1 - k_2\theta_i}{1 - \mathcal{E}^\circ(\theta)} \; > \; 0 \quad \forall \, \theta.$$

Hence, if $\frac{\partial e_i^\circ(\theta)}{\partial \theta_i} > 0$ for almost all $\theta_{-i} \in \Theta_{-i}$ then the first order condition for incentive compatibility will be violated.

To summarize, if e° is not compressed, then a country can move closer to its preferred extraction rate, not by cheating on the quota, but rather by misreporting its type. It will mimic a type that is allotted a higher extraction quota.

Naturally, this is a heuristic argument. It suggests that any optimal extraction plan must be compressed locally but does not rule out, for instance, step functions.

It nevertheless suggests a broader intuition. Namely, a free rider problem will exist with any solution to the planner's problem. Because of the "complementarity" between a country's actions and its reports, if any solution were non-compressed then a country can move closer to its preferred extraction rate by manipulating its report. The planner has no additional instrument to offset to this incentive to manipulate. The planner's only instrument is $\mathbf{e}$ and so full compression cannot be avoided.

5.2 Transferable Utility and Side Payments

In contrast with the logic above, environments with transferable utility provide to the planner an additional instrument — side payments. In the literature on self-enforcing environmental agreements (absent private information) the inclusion of monetary transfers or side payments in a climate agreement is often referred to as *issue linkage* (see Finus (2001)).

The transferable utility (TU) model provides a stark contrast with the full compression result. Using the expression for payoffs in (21) consider the following variant of the model: after receiving a reported profile, the planner can divide up the constant amount $nk_0 = n\frac{\log \omega_0}{1-\delta\gamma}$ and allocate it to the countries according to some formula. The allotment can include both positive (transfers) and negative values (taxes), depending on the reported profile. This allotment constitutes a transferable utility payment to/from each country.

Formally, let $s_i(\tilde{\theta})$ denote a tax/transfer from/to country i given the reported profile $\tilde{\theta}$. One can have either $s_i(\tilde{\theta}) > 0$ or $s_i(\tilde{\theta}) < 0$. A *balanced* transfer scheme satisfies $\sum_i s_i(\theta) = nk_0$ for every profile θ (where k_0 is defined in (22)). With this transfer, the country's payoff under truth-telling is

$$U_i \;=\; s_i(\theta) \;+\; k_1 \log(1-\mathcal{E}^\circ(\theta)) \;-\; \theta_i k_2 \log\left(\frac{1-\mathcal{E}^\circ(\theta)}{e_i^\circ(\theta)}\right). \tag{23}$$

Comparing (23) with the no-transfer payoff in Equation (21), we see that the constant k_0 is replaced with a type-contingent transfer $s_i(\theta)$.

Proposition 3 *There exists a PBE with a balanced transfer scheme $s = (s_1, \ldots, s_n)$ that implements the first-best (full information optimal) quota $c^*(\theta)$.*

Proposition 3 indicates that when the payoffs admit a TU representation, then the previous result is reversed and the first best is implementable.

Compression in the original model occurs precisely because each of the separable components of the payoffs are either independent of θ — as is the case with k_0 — or they jointly depend on θ via the extraction profile e°. By contrast, transferable utility imparts a degree of freedom to the planner by adding a linear type-contingent payoff that does not depend directly on e°.

The proof in the Appendix constructs a transfer scheme that adapts a well-known result of d'Aspremont and Gerard-Varet (1979) to the present model. In this construction, the transfer to country i is

$$
\begin{aligned}
s_i(\tilde{\theta}_i, \theta_{-i}) \;=\; & k_0 \;+\; \int_{\theta_{-i}} \sum_{j \neq i} [U_j(\omega_0, \mathbf{c}^{\circ}(\tilde{\theta}_i, \theta_{-i}), \theta_j)] \; dF_{-i} \\
& -\; \frac{1}{n-1} \sum_{j \neq i} \int_{\theta_{-j}} \sum_{k \neq j} [U_k(\omega_0, \mathbf{c}^{\circ}(\theta), \theta_k)] \; dF_{-j}
\end{aligned} \tag{24}
$$

In this construction, the average transfer is $k_0 = \frac{\log \omega_0}{1-\delta\gamma} > 0$. Hence, on average, countries receive the value k_0 of accumulated resource stock just as they did without the transfer. In relative terms, however, high type realizations are penalized, receiving a lower value than k_0. Meanwhile, low types receive a higher value than k_0. When aggregated over all countries, the transfers sum to nk_0.

In the Appendix proof, we verify that i's transfer depends only payoffs of other countries, and an aggregated part that integrates out i's type. By turning what previously was a constant term into a type contingent one, s_i aligns the payoff of i with those of the planner. Consequently, i has no incentive to lie. Note that the payments do not affect incentives at the compliance stage, and so the compliance arguments are unaffected.

The addition of transferable utility to our model is not simply a relaxation of an institutional constraint (i.e. "forbidding" side payments between countries), but is rather a fundamental change in the payoffs of each country's representative citizen. It is an open question whether and to what extent transfers can improve things under non-transferable utility (NTU) — as is the case in the present model. For a number of reasons, NTU transfers are problematic in our framework. First, the common NTU unit of account here is a resource unit such as carbon. The designer could conceivably propose an emissions trading system. Given the Revelation Principle, however, any trading mechanism that the designer could propose would be a special case of a quota system as defined already.[27]

Second, the direction of transfer indicated in (24) is regressive. To improve welfare, the scheme will at some point involve transfers from poor countries to rich ones. Specifically, in order to counteract compression, countries that realize shocks with high usage value (relative to its value of conservation) would be required to subsidize

[27]NTU transfers in other dimensions include non-climate benefits, technology transfers, trade terms, etc. To add these would require a substantially richer and higher dimensional model, but a worthwhile extension.

those with low usage value. Yet, the large projected increase in usage by developing countries, coupled with per capita reductions in GHG production by wealthier countries, means that relatively poorer countries will eventually be the ones with high usage value.[28]

Finally, the presence of side payments or any other form of issue linkage compounds transaction costs of reaching a global agreement. Along this line of argument Calcott and Petkov (2012) shows how cross-country heterogeneity reduces the possibility of an efficient implementation of transfers. They model the case of heterogeneous countries, under full information, focusing on transfers that are time invariant, linear in emissions, and consistent with budget balance. They find that heterogeneity reduces the scope for penalty schemes to jointly satisfy desirable emissions reduction.

5.3 Robustness Issues

Credibility. As with the compliance scheme in the full information environment, one can ask whether a planner will carry out the disclosure rule. At first glance, this may seem problematic. Consider, for instance, a truth-telling equilibrium that implements the optimal quota. Since the quota is fully compressed, each country has no incentive to lie about its type in date $t = 0$. This means that the IA now has full information about countries' types heading into date $t = 1$. Without the ability to commit to a mechanism at $t = 0$ the IA would make use of this information to implement the first-best solution in date $t = 1$ and beyond. However, this destroys the initial incentive for truthful disclosure.[29]

Consequently, if the IA can re-optimize at any date, then the equilibrium must be constructed so that disclosure strategies are pooling. Fortunately, this is easily done since the mechanism is already known to be fully compressed both on and off path.

Commons versus Market Incentives. One can also ask why the full compression holds in the commons environment but not necessarily in market frameworks. The countries in this environment can be compared to seller's in a procurement auction,

[28]See projected usage estimates from the Carbon Dioxide Information Analysis Center (CDIAC), *cdiac.ornl.gov*.

[29]This "ratchet effect" was originally observed by Roberts (1984) in a dynamic Mirleesian model of optimal taxation.

each attempting to bid for the right to emit carbon. A few key differences arise. In a standard auction or oligopoly there are trade offs in the production quota between seller types. An increase in the market share for a type θ_i, for instance, must be compensated by either a reduction in the share given to another seller-type on average. Given the hard trade offs in these cases, productive efficiency will usually require that the quota prescribe different levels for different seller types. For this reason, optimal mechanisms will not generally be compressed.[30]

One exception to the "un-compressed" optimal mechanism occurs when hazard rates F_i/f_i are strictly increasing. In that case, the loss from inefficient production is offset by allocative concerns of the planner. Athey and Bagwell (2008) for instance, analyze a repeated oligopoly setting in which firms receive serially correlated cost shocks each period. They show that the optimal production quota for colluding firms is fully compressed (which they refer to as "rigid") when either hazard rates are increasing or the maximum possible compensation from monopoly pricing is large enough.

McAfee and McMillan (1992) show a similar result when buyers collude in a static procurement auction under free entry. Finally, Lewis (1996) argues that compression would occur naturally when government uses standard policy instruments to regulate environmental externalities in a perfectly competitive static environment with privately observed costs and benefits of pollution. He shows, however, that without constraints on the instrument, the incentive-efficient mechanism is not compressed.

These results depend critically on the fact that market shares are bounded and in fact add up to one. This means that there are one-for-one trade offs across types in the oligopoly framework. By contrast, there is no such aggregate constraint in the global commons framework. The stored resource ω_t is non-rivalrous, and so all countries' values for, say, carbon storage can be increased simultaneously by simply withholding carbon consumption. For these reasons, the results from oligopolies (or procurement auctions) cannot be applied to the present model.

Imperfect Persistence. The compression result can be extended to a more general model of imperfect persistence. Specifically, a country is hit with a shock to its

[30]See Pavan, Segal, and Toikka (2012) for a comprehensive characterization of optimal mechanisms in Markov models of private information.

resource type each period according to a stationary Markov process. If that process is a martingale, then the quota will remain compressed.

Hence, two countries that start identically will typically evolve very different benefits and costs of resource use. They are, nevertheless, prescribed the same consumption quota throughout. The extension and result is described in an external appendix.[31]

The Parametric Structure. Finally, one can ask whether the results hold under a more general parametric setup. As a general matter, one can only speculate. Here we examine a generalization based on a utility function adapted from Antoniadou *et al.* (2013) (the "AKM model" from here on).[32] We extend the AKM model to a heterogeneous agent environment, and in an external appendix derive the Euler equations for the full information case.[33]

The law of motion in AKM model is given by

$$\omega_{t+1} = \left[\gamma(\omega_t - C_t)^{1-\frac{1}{\eta}} + (1-\gamma)\phi^{1-\frac{1}{\eta}}\right]^{\frac{\eta}{\eta-1}} \tag{25}$$

where $\eta > 0$. The larger is the second term in $[\cdot]$ the lower is the incremental effect of current resource depletion on future stock.

Now consider a flow utility for country i of

$$u_i = \theta_i \frac{C_{it}^{1-\frac{1}{\eta}} - 1}{1-\frac{1}{\eta}} + (1-\theta_i)\frac{(\omega_t - C_t)^{1-\frac{1}{\eta}} - 1}{1-\frac{1}{\eta}}. \tag{26}$$

The flow payoff in AKM model corresponds to Equation (26) in the special case where $\theta = 1$. Equation (26) simultaneously varies the elasticity of substitution η between extraction and preservation and the relative payoff weights given to each in u_i.

Notice that when taking the limits $\eta \to 1$, $\gamma \to 1$, and setting $\phi = 1$, Equations (25) and (26) converge to Equations (1) and (2), respectively. Hence this model is a generalization of our benchmark model.

[31]The argument is largely based on the proof of Proposition 2. It combines Part 1 of the original result, the martingale assumption which allows one to apply recursive logic. The proof of compliance mimics the steps of the proof of Proposition 2 (part 2). See faculty.georgetown.edu/lagunofr/optimal-resource10-External-Appendix2.pdf for the formulation and details.

[32]We thank a referee for suggesting the extension.

[33]The external appendix can be found on faculty.georgetown.edu/lagunofr/optimal-resource10-External-Appendix1.pdf.

For this more general model, the external appendix displays the Euler equations for the full information case.[34] We verify that the planner's problem is stationary, but does not admit an explicit solution. It turns out that the implementation in Proposition 1 holds when the CES parameter is in the set $(0, 1)$, but not when it is greater than 1. Basically, payoffs are bounded in extension when $\eta \geq 1$. But the construction of punishments in Proposition 1 applies only to commons problems in which payoffs are unbounded below. When payoffs are bounded, and an alternative construction of punishments, based on a Nash reversion strategy is worked out in Harrison and Lagunoff (2016). They show, however, the reversion strategies implements the optimum only for certain parameters.

It is unclear whether Proposition 2 can be extended to this case. Future research, making use of computational methods, would help to resolve this question.

6 Conclusion

This paper studies dynamic mechanisms for global commons with environmental externalities. Using carbon consumption as the leading example, we generalize the dynamic resource extraction game of Levhari and Mirman (1980) to allow for direct, heterogeneous benefits of resource conservation across countries. We examine the case where countries incur payoff/technology shocks. These shocks alter the way that countries evaluate the relative benefits and costs of carbon consumption over time.

An optimal quota system is an international agreement that assigns emissions restrictions to each country as a function of the sequence of realized type profiles such that it be implementable in PBE. The PBE builds in the idea of sequential self-enforcement in both compliance and disclosure.

Our main result is that the optimal quota system is fully compressed. The result is stark, as it suggests that the quota can only be tailored to ex ante differences between countries. Among other things, it should not vary with the realized evolution of a country's climate costs or its resource needs.

[34] These set of equations, beside helping to study the robustness of our results, can be helpful as a new baseline model with which to explore other issues in resource dynamics.

The results stand in contrast to with many actual international proposals (see Bodansky (2004) for a survey). Most of these advocate maximal flexibility in making adjustments particular characteristics of each country. Article 4 in the UNFCCC explicitly references the need to account for "the differences in these Parties' starting points and approaches, economic structures and resource bases, the need to maintain strong and sustainable economic growth, available technologies and other individual circumstances, as well as the need for equitable and appropriate contributions by each of these Parties to the global effort ..."

The results may be reconciled to these approaches to the extent that transfers can be used, or that some information about local shocks is globally observed. Indeed, our study suggests a critical role for sided payments, although not, as many policy makers suggest, for purposes of distributive justice (again see Bodansky (2004)). Both distributive concerns and political economy constraints may limit or negate the sorts of transfers or side payments needed to overcome the information frictions of the type we model.

As for increased information, inefficient compression could be mitigated if the data gathering process occurs above and beyond the reach of sovereign filters. This is an important consideration for the climate science itself. As before, any mitigating mechanism based on globally public information must be self enforcing at compliance stage or should be enforceable by the International agent. In our view, the disclosure of country-specific observations of economic costs and benefits remains problematic because these costs and benefits rely heavily on each country's yearly disclosure of its national income accounts. Public provision of "investigative" resources to generate such information, however, might overcome this problem. It seems an excellent topic for future study.

7 Appendix: Proofs of the Results

7.1 Derivation of Planner's Solution

.

Starting from the Euler equation

$$\frac{\theta_i}{e_{it}} - \sum_{j=1}^{n} \frac{(1-\theta_j)}{(1-\mathcal{E}_t)} + \sum_{j=1}^{n} \delta[\frac{\partial U_j(\omega_{t+1},\mathbf{e},\theta_j)}{\partial \omega_{t+1}} \frac{\partial \omega_{t+1}}{\partial e_{it}}] \;=\; 0$$

with $\frac{\partial \omega_{t+1}}{\partial e_{it}} = -\gamma \omega_t^{\gamma}(1-\mathcal{E}_t)^{\gamma-1}$, then the first order condition becomes

$$\frac{\theta_i}{e_{it}} - \sum_{j=1}^{n} \frac{(1-\theta_j)}{(1-\mathcal{E}_t)} \;=\; \gamma\delta \sum_{j=1}^{n} [\frac{\partial U_j(\omega_{t+1},\mathbf{e},\theta_{j\,t+1})}{\partial \omega_{t+1}} \omega_t^{\gamma}(1-\mathcal{E}_t)^{\gamma-1}] \tag{27}$$

Differentiating the value function $U_j(\omega_{t+1},\mathbf{e},\theta_j)$ with respect to ω_{t+1} yields,

$$\frac{\partial U_j(\omega_{t+1},\mathbf{e},\theta_j)}{\partial \omega_{t+1}} \;=\; \frac{1}{\omega_{t+1}} \;+\; \delta\gamma \frac{\partial U_j(\omega_{t+2},\mathbf{e},\theta_j)}{\partial \omega_{t+2}} \omega_{t+1}^{\gamma-1}(1-\mathcal{E}_{t+1})^{\gamma}$$

Substituting this into the first order condition, we obtain

$$\begin{aligned}
& \frac{\theta_i}{e_{it}} - \sum_{j=1}^{n} \frac{(1-\theta_j)}{(1-\mathcal{E}_t)} \\
= \;& \delta\gamma \sum_{j=1}^{n} \left\{ \frac{1}{\omega_{t+1}} \;+\; \frac{\partial U_j(\omega_{t+2},\mathbf{e},\theta_j)}{\partial \omega_{t+2}} \omega_{t+1}^{\gamma-1}(1-\mathcal{E}_{t+1})^{\gamma} \omega_t^{\gamma}(1-\mathcal{E}_t)^{\gamma-1} \right\} \\
= \;& \frac{1}{1-\mathcal{E}_t}\gamma\delta \sum_{j=1}^{n} \left\{ \frac{\omega_t^{\gamma}(1-\mathcal{E}_t)^{\gamma}}{\omega_{t+1}} + \frac{\partial U_j(\omega_{t+2},\mathbf{e},\theta_j)}{\partial \omega_{t+2}} \omega_{t+1}^{\gamma-1}(1-\mathcal{E}_{t+1})^{\gamma} \omega_t^{\gamma}(1-\mathcal{E}_t)^{\gamma} \right\} \\
= \;& \frac{1}{1-\mathcal{E}_t} \left(n\delta\gamma + \delta\gamma \sum_{j=1}^{n} \frac{\partial U_j(\omega_{t+2},\mathbf{e},\theta_j)}{\partial \omega_{t+2}} \omega_{t+1}^{\gamma}(1-\mathcal{E}_{t+1})^{\gamma} \right) \\
= \;& \frac{1}{1-\mathcal{E}_t} \left(n\delta\gamma + \delta\gamma \sum_{j=1}^{n} [(\frac{\theta_j(1-\mathcal{E}_{t+1})}{e_{j\,t+1}} - \sum_{j=1}^{n}(1-\theta_j))] \right)
\end{aligned}$$

Reorganizing terms, we obtain

$$\frac{\theta_i(1-\mathcal{E}_t)}{e_{it}} - \sum_{j=1}^{n}(1-\theta_j) \;=\; n\delta\gamma + \delta\gamma \left[\frac{\theta_i(1-\mathcal{E}_{t+1})}{e_{i\,t+1}} - \sum_{j=1}^{n}(1-\theta_j) \right] \tag{28}$$

which is Equation (7) in the main body of the paper.

Define $s_{it} \;\equiv\; \frac{\theta_i(1-\mathcal{E}_t)}{e_{it}} - \sum_{j=1}^{n}(1-\theta_{jt})$. Then

$$s_{it}(\theta_i) = n\delta\gamma + \delta\gamma s_{i\,t+1}(\theta_i).$$

Forward iteration yields a steady state $s_{it} = \frac{n\delta\gamma}{(1-\delta\gamma)}$ and plugging in the last equation we obtain

$$\frac{\theta_i(1-\mathcal{E}_t)}{e_{it}} - \sum_{j=1}^{n}(1-\theta_j) \;=\; \frac{n\delta\gamma}{(1-\delta\gamma)}.$$

Then:

$$
\begin{aligned}
\frac{\theta_i(1-\mathcal{E}_t)}{e_{it}} - \sum_{j=1}^{n}(1-\theta_j) &= \frac{n\delta\gamma}{(1-\delta\gamma)} && (29)\\
\theta_i(1-\mathcal{E}_t)(1-\delta\gamma) - e_{it}\sum_j(1-\theta_j)(1-\gamma\delta) &= e_{it}\gamma\delta n \\
\theta_i(1-\mathcal{E}_t)(1-\gamma\delta) - e_{it}(n-\sum_j\theta_j)(1-\gamma\delta) &= e_{it}\gamma\delta n && (30)\\
\sum_j\theta_j(1-\mathcal{E}_t)(1-\gamma\delta) - (1-\gamma\delta)\mathcal{E}_t(n-\sum_j\theta_j) - \mathcal{E}_t\gamma\delta n &= 0 && (31)\\
\sum_j\theta_j - n\mathcal{E}_t - \sum_j\theta_j\gamma\delta &= 0 && (32)
\end{aligned}
$$

using Equations (29)-(32) we get

$$e_i^*(\theta_i) = \frac{\theta_i(1-\delta\gamma)}{n} \tag{33}$$

which is equivalent to Equation (9) in the main text.

7.2 Proof of Proposition 1

Fix a profile θ. As we have already shown that $\mathbf{c}^*(\theta)$ maximizes $\sum_i U_i$ without the equilibrium constraint, it remains to show that $\mathbf{c}^*(\theta)$ can be implemented by a SPE. For ease of exposition, we will work with extraction rates e_i rather than levels c_i.

To prove the result, we will construct a sequence $\{e^\tau(\theta)\}_{\tau=0}^{\infty}$ of extraction profiles, and a strategy profile σ satisfying for each i:

(i) $\sigma_i(h^0,\theta) = \omega_0 e_i^0(\theta)$,

(ii) for each $\tau \geq 0$ and each $t \geq \tau$, if $h^t(e^\tau(\theta))$ represents history in which some country deviated unilaterally from profile $e^\tau(\theta)$ in date $t-1$, then set

$$\sigma_i(h^t(e^\tau(\theta)), \theta) = \omega_t e_i^{\tau+1}(\theta)$$

(i) For all other histories h^t let $\sigma_i(h^t, \theta) = \omega_0 e_i^0(\theta)$.

Our task is to construct the sequence $\{e^\tau(\theta)\}_{\tau=0}^{\infty}$ in such a way that σ is a SPE, i.e., $V_i(h^t, \sigma, \theta) \geq V_i(h^t, \hat{\sigma}_i, \sigma_{-i}, \theta)$ for all i, all $\hat{\sigma}_i$, and all histories h^t, and along the equilibrium path h^{*t} at date t, $\sigma_i(h^{*t}, \theta) = c_{it}^*(\theta)$.

Our sequence is constructed recursively as follows. Starting from ω_0, for any stationary dynamic path $\mathbf{e}$ of usage rates, the long run payoff to a country i can be expressed as

$$(34)\qquad \begin{aligned} &\frac{\log(\omega_0)}{1-\delta\gamma} + \sum_{t=0}^{\infty} \delta^t \left[\left(\frac{1}{1-\delta\gamma} - \theta_i \right) \log(1-\mathcal{E}) + \theta_i \log e_i \right] \\ \equiv\ &\frac{\log(\omega_0)}{1-\delta\gamma} + \sum_{t=0}^{\infty} \delta^t \, u_i \end{aligned}$$

Here, u_i captures the long run effect on payoffs of the stationary extraction profile $\mathbf{e}$ chosen at each date t. We refer to u_i as the *flow payoff* even though it includes future as well as present effects of the current profile $\mathbf{e}$ through its effect on the resource stock ω_t. The critical feature used in the proof is the fact that each "flow" payoff u_i is unbounded below. In the rest of the proof we make use of this notation and, moreover, drop the first term $\frac{\log(\omega_0)}{1-\delta\gamma}$ which will cancel in any comparison with an alternative long run payoff.

Let $\mathbf{e}^*(\theta)$ denote the corresponding path of usage rates in the optimal quota $\mathbf{c}^*(\theta)$. Recall that $e_i^*(\theta) = \frac{\theta_i(1-\delta\gamma)}{n}$. Since $\mathbf{e}^*(\theta)$ is stationary, i.e., $\mathbf{e}_t^*(\theta) = \mathbf{e}_{t'}^*(\theta)$ for any pair of dates t and t', it yields a payoff

$$(35)\qquad V_i^*(\theta) \equiv \frac{1}{1-\delta} \left[\left(\frac{1}{1-\delta\gamma} - \theta_i \right) \log(1-\mathcal{E}^*(\theta)) + \theta_i \log e_i^*(\theta) \right]$$

to each country i (and dropping the term $\frac{\log(\omega_0)}{1-\delta\gamma}$).

Working with rates rather than levels, we construct a recursive sequence of usage profiles $\{e^\tau(\theta)\}_{\tau=0}^{\infty}$ as follows. Let $\mathbf{e}^0(\theta) = \mathbf{e}^*(\theta)$ and $V_i^0(\theta) = V_i^*(\theta)$.

Next, for $\tau \geq 1$, let $e^{\tau-1}(\theta)$ be a profile of extraction rates. The profile $e^{\tau}(\theta)$ is defined as follows.

For each τ, all countries choose $e^{\tau}(\theta)$ (to be defined shortly) for that period. This yields each country a flow payoff of $u_i^{\tau}(\theta)$. After one period, the countries revert to optimal extraction rates, $\mathbf{e}^0(\theta)$, if there is no subsequent deviation.[35] The payoff $V_i^{\tau}(\theta)$ is therefore defined by

$$V_i^{\tau}(\theta) \;\equiv\; u_i^{\tau}(\theta) + \delta V_i^0(\theta)$$

Using the definition in (34) of an arbitrary flow payoff, the payoff in state τ is given by

$$(36) \qquad V_i^{\tau}(\theta) = \left(\frac{1}{1-\delta\gamma} - \theta_i\right) \log(1 - \mathcal{E}^{\tau}(\theta)) \;+\; \theta_i \log e_i^{\tau}(\theta) \;+\; \delta\, V_i^0(\theta)$$

To complete the recursive definition, we need to define $e^{\tau}(\theta)$. This will be constructed to satisfy the incentive constraint in stage $\tau - 1$. Specifically, if it turns out that a country deviates in state $\tau - 1$ in some period then the countries transition to the state τ in the following period so that extraction rates in the next period are given by $e^{\tau}(\theta)$. Consequently, for each $\tau \geq 1$, $e^{\tau}(\theta)$ is defined to satisfy:

$$(37) \qquad V_i^{\tau-1}(\theta) \;\geq\; \bar{u}_i^{\tau-1}(\theta) \;+\; \delta V_i^{\tau}(\theta)$$

where $\bar{u}_i^{\tau-1}(\theta) \;=\; \max_{e_i} \left[\left(\frac{1}{1-\delta\gamma} - \theta_i\right) \log(1 - \mathcal{E}_{-i}^{\tau-1}(\theta) + e_i) \;+\; \theta_i \log e_i\right]$ is i's flow payoff if it chooses a best response to $e^{\tau-1}(\theta)$ in the current period.

Using the definition of V_i^{τ} in (36), one can show that the incentive constraint (37) holds if for all i,

(38)
$$\bar{u}_i^{\tau-1}(\theta) + \delta(1-\delta)V_i^0(\theta) \;\geq\; \bar{u}_i^{\tau-1}(\theta) + \delta \left[\left(\frac{1}{1-\delta\gamma} - \theta_i\right) \log(1 - \mathcal{E}^{\tau}(\theta)) \;+\; \theta_i \log e_i^{\tau}(\theta)\right]$$

Clearly, these inequalities (one for each country) can always be made to hold by choosing $\mathcal{E}^{\tau}$ sufficiently close to one. Note that it will necessarily be the case that $\mathcal{E}^{\tau} > \mathcal{E}^{\tau-1}$.

To summarize, the sequence $\{e^{\tau}(\theta)\}_{\tau=0}^{\infty}$ is recursively constructed so that for each $e^{\tau-1}(\theta)$, $e^{\tau}(\theta)$ is chosen to satisfy these incentive constraints. Then σ is constructed to

[35] We thank a referee for suggesting this simplified construction.

give precisely $e^{\tau}(\theta)$ in any period in which a unilateral deviation occurred in $e^{\tau-1}(\theta)$ after which point the strategy takes the countries back to $e^0(\theta) = e^*(\theta)$. Multilateral deviations revert to $e^0(\theta)$ as well. By construction, the profile is subgame perfect and implements $\mathbf{c}^*(\theta)$. ∎

7.3 Proof of Proposition 2

7.3.1 A Basic Outline

Before proceeding, we give a more formal outline of the argument sketched in Section 5.1.

Recalling the discussion in Serction 5.1, the proof of Part 1 proceeds by first introducing a "relaxed planner's problem" as shown in Equation (20). Since the relaxed problem ignores off-path deviations, all the subsequent arguments here refer to the on-path incentive constraints in (19).

Without the incentive constraint (19), the solution to the relaxed problem in (20) coincides the with full information extraction plan (Lemma 1 below). With incentive constraints, the solution to the relaxed problem can be reformulated using standard Envelope and monotonicity conditions from the theory of optimal auctions (Lemma 2 below).

Both Lemma 1 and Lemma 2 are used in Lemmata 3-6 (in the proof) to rule out all possible forms of non-compression. The idea is as follows. For any profile θ_{-i} of others' types, the full information optimal extraction rate for country i possessing the highest type $\overline{\theta}$ is

$$e_i^*(\overline{\theta}, \theta_{-i}) \; = \; \overline{e} \; \equiv \; \overline{\theta}(1 - \delta\gamma)/n.$$

This extraction rate serves as a cutoff. Any solution to the indirect problem must be consistent with an extraction plan $e_i^\circ(\theta)$ for each i that lies *below* $\overline{e}$. This is shown below in Lemma 3.

By contrast, under the Dispersion restriction, if a type could freely choose its extraction rate e_i then it would choose a larger extraction rate than $\overline{e}$. This is shown in Lemma 4.

Hence, if e° is not compressed, then a country can move closer to its preferred extraction rate, not by cheating on the quota, but rather by misreporting its type. It will mimic a type that is allotted a higher extraction quota. This is shown in Lemma 5.

To prove Part 2 of the Proposition, we need to show that the relaxed-optimal quota can be implemented by a full-blown PBE in which a country's prescribed usage rate after any history is stationary and compressed.

To implement the quota $\mathbf{e}^{\circ}$ one might think that the IA could simply use the sanctions constructed in the full information case. That logic (outlined in the proof of Proposition 1) requires credible punishment at each counterfactual stage depends on even harsher punishment if the countries fail to carry out the sanction.

When shocks are private, that construction needs modification. In any PBE with truth-telling at the disclosure stage, the IA takes at face value any disclosed profile, manipulated or otherwise. Consequently, a country may have an incentive to report a type consistent with the least punitive punishment which, in turn, may give it the incentive to violate the quota itself. Notice that this problem arises whenever the prescribed usage along the punishment path varies across type. So, even though the quota $\mathbf{c}^{\circ}$ itself is compressed, it may be difficult to implement in Perfect Bayesian equilibria if the sanctions are not. The result in Part 2 by-passes this problem by showing that the optimal quota can be implemented by compressed sanctions.

7.3.2 The Proof

Part 1. We begin by establishing some facts about the planner's relaxed problem in (20).

Lemma 1 *The full information solution* $e_i^*(\theta_i) \equiv \theta_i(1-\delta\gamma)/n$ *for country* i *solves (20) if the truth-telling constraint (19) is dropped.*

Proof of Lemma 1. By dropping the constraint (19), truthful revelation is assumed to hold at $t = 0$, and so the planner implements the full information optimum $\mathbf{e}^*$ after receiving the profile θ. The Lemma therefore follows. ∎

Call an extraction plan $\mathbf{c}^\circ$ *symmetric* if for any ω and any permutation $\nu : \{1,\dots,n\} \to \{1,\dots,n\}$ on the index set of countries and any profile θ with θ_ν corresponding to the permutation of θ under ν, we have,

$$c_i^\circ(\omega,\theta) = c_j^\circ(\omega,\theta_\nu) \quad \text{whenever} \quad j = \nu(i).$$

The solution to the planner's relaxed problem (20) is symmetric given the symmetry of payoffs, incentive constraints, and welfare weights in the planner's problem.

Let $\mathbf{c}^\circ$ denote a candidate symmetric solution to (20), and let $\mathbf{e}^\circ$ denote the rates corresponding to $\mathbf{c}^\circ$. For any equilibrium (σ,μ) which implements $\mathbf{c}^\circ$, if (σ,μ) satisfies the truth-telling constraint (19), then any arbitrary report $\widetilde{\theta}_j$ by country j should be believed to be j's true type with probability one by country i. By the Revelation Principle, for any PBE one can always construct a payoff-equivalent PBE with truth-telling strategies and so we limit our attention to these.

Consequently, if profile $\widetilde{\theta}$ is disclosed by all the countries, then a country of type θ_i has a long- run payoff under a quota (in rates) $\mathbf{e}^\circ$ given by,

(39)

$$\begin{aligned}
U_i(\omega_0,\mathbf{e}^\circ(\tilde{\theta}),\theta_i) &= \sum_{t=0}^{\infty}\delta^t\left[\theta_i\log\omega_t e_{it}^\circ(\widetilde{\theta}) + (1-\theta_i)\log\omega_t(1-\mathcal{E}_t^\circ(\widetilde{\theta}))\right]\\
&= \sum_{t=0}^{\infty}\delta^t\left[\log\omega_t + \theta_i\log e_{it}^\circ(\widetilde{\theta}) + (1-\theta_i)\log(1-\mathcal{E}_t^\circ(\widetilde{\theta}))\right]\\
&= \sum_{t=0}^{\infty}\delta^t\log\left(\omega_0^{\gamma^t}\prod_{s=0}^{t-1}(1-\mathcal{E}_s^\circ(\widetilde{\theta}))^{\gamma^{t-s}}\right) \;+\; \theta_i\sum_{t=0}^{\infty}\delta^t\log e_{it}^\circ(\widetilde{\theta})\\
&\qquad +\; (1-\theta_i)\sum_{t=0}^{\infty}\delta^t\log(1-\mathcal{E}_t^\circ(\widetilde{\theta}))\\
&= \frac{\log\omega_0}{1-\delta\gamma} \;+\; \sum_{t=0}^{\infty}\delta^t\sum_{s=0}^{t-1}\gamma^{t-s}\log(1-\mathcal{E}_s^\circ(\widetilde{\theta})) \;+\; \sum_{t=0}^{\infty}\delta^t\log(1-\mathcal{E}_t^\circ(\widetilde{\theta}))\\
&\qquad -\; \theta_i\sum_{t=0}^{\infty}\delta^t\log\left(\frac{1-\mathcal{E}_t^\circ(\widetilde{\theta})}{e_{it}^\circ(\widetilde{\theta})}\right)\\
&= \frac{\log\omega_0}{1-\delta\gamma} + \sum_{t=0}^{\infty}\delta^t\sum_{s=0}^{t}\gamma^{t-s}\log(1-\mathcal{E}_s^\circ(\widetilde{\theta})) \;-\; \theta_i\sum_{t=0}^{\infty}\delta^t\log\left(\frac{1-\mathcal{E}_t^\circ(\widetilde{\theta})}{e_{it}^\circ(\widetilde{\theta})}\right)
\end{aligned}$$

Observe that $U_i(\omega_0,\mathbf{e}^\circ(\tilde{\theta}),\theta_i)$ is a strictly concave, hence single peaked, function of $\{e_t(\tilde{\theta})\}$. Multiplying U_i by the standard normalizing constant $(1-\delta)$, the time t weights sum to one. Consequently there exists a "certainty equivalent" extraction plan, i.e., a plan $\mathbf{e}$ satisfying $e_{i\,t}(\tilde{\theta}) = e_{i\,t'}(\tilde{\theta})$ for all i, t, t', and $\tilde{\theta}$ (time stationarity), and satisfying $U_i(\omega_0, \mathbf{e}(\widetilde{\theta}), \theta_i) = U_i(\omega_0, \mathbf{e}^{\circ}(\widetilde{\theta}), \theta_i)$ for all i and all θ.[36] By symmetry of $\mathbf{e}$, the equivalence holds for each country $i = 1, \ldots, n$. It also holds holds profile by profile, and so $\mathbf{e}$ satisfies the incentive constraints.

Consequently we restrict attention to stationary extraction plans without loss of generality. For any extraction profile e, define

$$r(e) \equiv \frac{\log \omega_0}{1-\delta\gamma} + \log(1-\mathcal{E}) \sum_{t=0}^{\infty} \delta^t \sum_{s=0}^{t} \gamma^{t-s} \quad \textit{and} \tag{40}$$

$$q_i(e) \equiv \frac{1}{1-\delta} \log\left(\frac{1-\mathcal{E}}{e}\right). \tag{41}$$

Then a stationary extraction profile $\mathbf{e}^{\circ}$ satisfies

$$U_i(\omega_0, \mathbf{e}^{\circ}(\tilde{\theta}), \theta_i) = r(e^{\circ}(\widetilde{\theta})) - \theta_i \, q_i(e^{\circ}(\widetilde{\theta})).$$

Reassuringly, this payoff is consistent with the payoff described in Equation (21).

This payoff has a simple interpretation. For any type profile $\widetilde{\theta}$ reported to the IA, r is the long run benefit of resource conservation. By contrast, q_i is the long run cost of conservation. This cost is the value of the foregone usage relative to one's actual usage. It is therefore naturally decreasing in one's actual usage. Notice that both r and q_i are strictly concave functions of the profile $e^{\circ}(\widetilde{\theta})$ for any $\tilde{\theta}$.

In what follows, we ignore the constant term in r. The interim values of r and q_i are given by

$$R_i^{\circ}(\theta_i) \equiv \int_{\theta_{-i}} r(e^{\circ}(\theta)) dF_{-i}(\theta_{-i}) \quad \textit{and} \quad Q_i^{\circ}(\theta_i) \equiv \int_{\theta_{-i}} q_i(e^{\circ}(\theta)) dF_{-i}(\theta_{-i}) \tag{42}$$

Hence, by definition,

$$\int_{\theta_{-i}} U_i(\omega_0, \mathbf{c}^{\circ}(\theta), \theta_i) dF_{-i}(\theta_{-i}) \equiv R_i^{\circ}(\theta_i) - \theta_i \, Q_i^{\circ}(\theta_i), \tag{43}$$

and the truth-telling constraint (19) may be expressed as

$$R_i^{\circ}(\theta_i) - \theta_i \, Q_i^{\circ}(\theta_i) \geq R_i^{\circ}(\tilde{\theta}_i) - \theta_i \, Q_i^{\circ}(\tilde{\theta}_i) \quad \forall\, \theta_i \;\; \forall\, \tilde{\theta}_i \tag{44}$$

[36] Because the discount weights can be interpreted as probability weights, the extraction plan $\mathbf{e}$ constitutes a certainty equivalent to the "random" (i..e, time varying) extraction plan $\mathbf{e}^{\circ}$.

Notice that (43) and (44) resemble payoffs in a standard procurement auction or oligopoly problem, though we emphasize that R° and Q° are not independent choices, but instead depend jointly on e°.

In general, for any pair of functions R_i and Q_i comprising the payoff $R_i(\theta_i) - \theta_i\, Q(\theta_i)$, standard arguments from optimal auctions (Myerson (1981)) show that any R_i and Q_i will satisfy the incentive constraint (44) only if Q_i is weakly decreasing and

$$R_i(\theta_i) \;-\; \theta_i\; Q_i(\theta_i) \;=\; R_i(\overline{\theta}) - \overline{\theta}\; Q_i(\overline{\theta}) + \int_{\tilde{\theta}_i=\theta_i}^{\overline{\theta}} Q_i(\tilde{\theta}_i)d\tilde{\theta}_i \tag{45}$$

Equation (45) is the standard envelope condition where the IC constraint for the highest type $\overline{\theta}$ does not bind. An integration by parts argument applied to (45) yields

$$\int_{\theta_i=\underline{\theta}}^{\overline{\theta}}[R_i(\theta_i) \;-\; \theta_i\; Q_i(\theta_i)]dF_i(\theta_i) \;=\; R_i(\overline{\theta}) - \overline{\theta}\; Q_i(\overline{\theta}) + \int_{\underline{\theta}}^{\overline{\theta}} F_i(\theta_i)Q_i(\theta_i)d\theta_i$$

We have thus shown that the planner's problem can be formulated as follows:

Lemma 2 *Any solution* $\mathbf{e}^\circ$ *to the planner's relaxed problem solves*

$$\begin{aligned} &\max_{\mathbf{e}^\circ} \;\sum_i \;\left[R_i^\circ(\overline{\theta}) - \overline{\theta}\; Q_i^\circ(\overline{\theta}) + \int_{\underline{\theta}}^{\overline{\theta}} F_i(\theta_i)Q_i^\circ(\theta_i)d\theta_i\right] \;= \\ &\max_{\mathbf{e}^\circ} \;\sum_i \;\left[\int_{\theta_{-i}} [r(e^\circ(\overline{\theta},\theta_{-i})) - \overline{\theta}\; q_i(e^\circ(\overline{\theta},\theta_{-i}))\;]dF_{-i}(\theta_{-i})\right. \\ &\qquad\qquad \left. + \int_{\underline{\theta}}^{\overline{\theta}} F_i(\theta_i) \int_{\theta_{-i}} q_i(e^\circ(\theta)dF_{-i}(\theta_{-i})d\theta_i \right] \end{aligned} \tag{46}$$

subject to Q_i° *weakly decreasing in* θ_i *for all* i.

We will refer to (46) as the *indirect problem.* We want to show that any extraction plan $\mathbf{e}^\circ$ that solves the indirect problem in (46), is fully compressed.

In the following Lemmata, $\mathbf{e}^\circ$ will denote a solution to the indirect problem in (46).

Lemma 3 *For every country i,*

$$e_i^{\circ}(\theta) \;\leq\; \bar{e}_i \equiv \overline{\theta}(1-\delta\gamma)/n \qquad a.e.\ \theta.$$

Proof of Lemma 3. Suppose by contradiction that $e_i^{\circ}(\theta) > \bar{e}_i$ for some country i on a set Θ^i of positive Borel measure.

Step 1. Observe first that by Lemma 1, $\bar{e}_i = e^*(\bar{\theta}, \theta_{-i})$. That is, $\bar{e}_i$ is the unconstrained socially optimal extraction rate assigned to country i when i is of type $\overline{\theta}$, regardless of the extraction types of other countries.

Hence, for *every* θ,

$$\sum_i \left[r(\bar{e}) - \overline{\theta}\ q_i(\bar{e})\ \right] \;>\; \sum_i \left[r(e^{\circ}(\overline{\theta}, \theta_{-i})) - \overline{\theta}\ q_i(e^{\circ}(\overline{\theta}, \theta_{-i}))\ \right] \tag{47}$$

Since

$$\int_{\theta}\left[\sum_i \left[r(e^{\circ}(\overline{\theta}, \theta_{-i})) - \overline{\theta}\ q_i(e^{\circ}(\overline{\theta}, \theta_{-i}))\ \right]\right] dF$$
$$= \sum_i \left[\int_{\theta_{-i}} [r(e^{\circ}(\overline{\theta}, \theta_{-i})) - \overline{\theta}\ q_i(e^{\circ}(\overline{\theta}, \theta_{-i}))\ ] dF_{-i}(\theta_{-i})\right]$$

it follows that

$$\sum_i \left[\overline{R}_i - \overline{\theta}\ \overline{Q}_i\right] \;>\; \sum_i \left[R_i^{\circ}(\overline{\theta}) - \overline{\theta}\ Q_i^{\circ}(\overline{\theta})\right] \tag{48}$$

Step 2. Define an extraction plan $\mathbf{e}^{**}$ by: for all $j \neq i$,

$$e_j^{**}(\theta) = \begin{cases} \bar{e}_j & if \quad \theta = (\bar{\theta}, \theta_{-k}), \text{ for some } k = 1, \ldots, n \\ e_j^{\circ}(\theta) & if \quad otherwise \end{cases}$$

That is, each country $j \neq i$ has same extraction profile as e° except when some country (possibly different from j) has the highest type $\bar{\theta}$. Next, define i's extraction profile by

$$e_i^{**}(\theta) = \begin{cases} \bar{e}_i & if \quad \theta \in \bigcup_{\theta_{-i}} \{(\bar{\theta}, \theta_{-i})\} \cup \Theta^i \\ e_i^{\circ}(\theta) & if \quad otherwise \end{cases}$$

As constructed, $e_i^{**}(\theta) \leq e_i^\circ(\theta)$ except possibly when $\theta_i = \bar{\theta}$.

The monotonicity of q_i in e implies

$$
\begin{aligned}
(49) \qquad & \int_{\underline{\theta}}^{\overline{\theta}} F_i(\theta_i) Q_i^{**}(\theta_i) d\theta_i \;\equiv\; \int_{\underline{\theta}}^{\overline{\theta}} F_i(\theta_i) \int_{\theta_{-i}} q_i(e^{**}(\theta)) dF_{-i}(\theta_{-i}) d\theta_i \\
\geq \; & \int_{\underline{\theta}}^{\overline{\theta}} F_i(\theta_i) \int_{\theta_{-i}} q_i(e^\circ(\theta)) dF_{-i}(\theta_{-i}) d\theta_i \;\equiv\; \int_{\underline{\theta}}^{\overline{\theta}} F_i(\theta_i) Q_i^\circ(\theta_i) d\theta_i
\end{aligned}
$$

Note that type profiles of the form $(\bar{\theta}, \theta_{-k})$ are of measure zero and hence are irrelevant for this inequality.

Step 3. Combining Equations (47) and (48) from Step 1 and Equation (49) from Step 2,

$$
\begin{aligned}
& \sum_i \left[R_i^{**}(\overline{\theta}) - \overline{\theta}\; Q_i^{**}(\overline{\theta}) + \int_{\underline{\theta}}^{\overline{\theta}} F_i(\theta_i) Q_i^{**}(\theta_i) d\theta_i \right] \quad \equiv \\
& \sum_i \left[\int_{\theta_{-i}} [r(e^{**}(\overline{\theta}, \theta_{-i})) - \overline{\theta}\; q_i(e^{**}(\overline{\theta}, \theta_{-i}))\;] dF_{-i}(\theta_{-i}) \right. \\
& \qquad \left. + \int_{\underline{\theta}}^{\overline{\theta}} F_i(\theta_i) \int_{\theta_{-i}} q_i(e^{**}(\theta)) dF_{-i}(\theta_{-i}) d\theta_i \right] \quad = \\
& \sum_i \left[\int_{\theta_{-i}} [r(\bar{e}) - \overline{\theta}\; q_i(\bar{e})\;] dF_{-i}(\theta_{-i}) \right. \\
& \qquad \left. + \int_{\underline{\theta}}^{\overline{\theta}} F_i(\theta_i) \int_{\theta_{-i}} q_i(e^{**}(\theta)) dF_{-i}(\theta_{-i}) d\theta_i \right] \quad > \\
& \sum_i \left[\int_{\theta_{-i}} [r(e^\circ(\overline{\theta}, \theta_{-i})) - \overline{\theta}\; q_i(e^\circ(\overline{\theta}, \theta_{-i}))\;] dF_{-i}(\theta_{-i}) \right. \\
& \qquad \left. + \int_{\underline{\theta}}^{\overline{\theta}} F_i(\theta_i) \int_{\theta_{-i}} q_i(e^\circ(\theta)) dF_{-i}(\theta_{-i}) d\theta_i \right] \quad \equiv \\
& \sum_i \left[R_i^\circ(\overline{\theta}) - \overline{\theta}\; Q_i^\circ(\overline{\theta}) + \int_{\underline{\theta}}^{\overline{\theta}} F_i(\theta_i) Q_i^\circ(\theta_i) d\theta_i \right],
\end{aligned}
$$

a contradiction of the optimality of e°. ∎

Lemma 4 *Fixing ω_0, define*

$$
BR_i(\theta_i, e_{-i}, F_{-i}) \;=\; \arg\max_{e_i} \int_{\theta_{-i}} U_i(\omega_0, \mathbf{e}(\theta), \theta_i) dF_{-i}
$$

Then $BR_i(\theta_i, e^\circ_{-i}, F_{-i}) > \bar{e}$ *for every type* θ_i*.*

Proof of Lemma 4. Fixing ω_0, define

$$\overline{BR}_i(\theta_i, e_{-i}(\theta_{-i})) \quad = \quad \arg\max_{e_i} \; U_i(\omega_0, \mathbf{e}(\theta), \theta_i)$$

Extraction rates $\overline{BR}_i(\theta_i, e_{-i}(\theta_{-i}))$ and $BR_i(\theta_i, e^\circ_{-i}, F_{-i})$ are country i's ex post and interim best response functions, respectively, when the initial state is ω_0. Indeed, one can think of $\overline{BR}_i(\theta_i, e_{-i}(\theta_{-i}))$ as corresponding to $BR_i(\theta_i, e^\circ_{-i}, F_{-i})$ when F_{-i} is a degenerate distribution that places full mass on θ_{-i}.

Since $U_i(\omega_0, \mathbf{e}(\theta), \theta_i)$ and $\int_{\theta_{-i}} U_i(\omega_0, \mathbf{e}(\theta), \theta_i) dF_{-i}$ are strictly concave functions of e_i, it follows that $\overline{BR}_i(\theta_i, e_{-i}(\theta_{-i}))$ and $BR_i(\theta_i, e_{-i}, F_{-i})$ are single valued.

Define

$$\theta^{min}_{-i} \in \arg\min_{\theta_{-i}} \overline{BR}_i(\theta_i, e_{-i}(\theta_{-i}))$$

(Note that θ^{min}_{-i} minimizes i's best response, but need not be, itself, a minimum on $[\underline{\theta}, \overline{\theta}]^{n-1}$.)

From (12),

$$\overline{BR}_i(\theta_i, e^\circ_{-i}(\theta^{min}_{-i})) = \theta_i(1 - \mathcal{E}^\circ_{-i}(\theta_i, \theta^{min}_{-i}))(1 - \delta\gamma)$$

Hence, a country's preferred extraction rate is decreasing in the anticipated extraction of others, and so θ^{min}_{-i} must maximize the total extraction $\mathcal{E}^\circ_{-i}(\theta_{-i})$ of other countries. Consequently,

$$BR_i(\theta_i, e^\circ_{-i}, F_{-i}) \;\geq\; \overline{BR}_i(\theta_i, e^\circ_{-i}(\theta^{min}_{-i})\;) \tag{50}$$

Hence, for any θ_i we have

$$
\begin{aligned}
BR_i(\theta_i, e^\circ_{-i}, F_{-i}) &\geq \overline{BR}_i(\theta_i, e^\circ_{-i}, \theta^{min}_{-i}) \\
&= \theta_i(1 - \mathcal{E}^\circ_{-i}(\theta_i, \theta^{min}_{-i}))(1 - \delta\gamma) \\
&\geq \theta_i(1 - \sum_{j \neq i} \overline{e}_j)(1 - \delta\gamma) \\
&\equiv \theta_i(1 - \tfrac{(n-1)}{n}\overline{\theta}(1 - \delta\gamma))(1 - \delta\gamma) \\
&> \frac{\overline{\theta}(1 - \delta\gamma)}{n} \\
&\equiv \bar{e}_i
\end{aligned} \tag{51}
$$

In this string, the first inequality follows from (50), the second inequality follows from Lemma 3. The third inequality follows from the Disperson restriction

$$
\underline{\theta} > \frac{1}{1 + (n-1)\delta\gamma} \tag{52}
$$

assumed in the Proposition. To verify this final inequality, observe that

$$
\begin{aligned}
&\underline{\theta} \; > \; \tfrac{1}{1+(n-1)\delta\gamma} \implies \\
&\underline{\theta} \; > \; \tfrac{1}{n-(n-1)(1-\delta\gamma)} \implies \\
&n - (n-1)(1-\delta\gamma)) \; > \; 1/\underline{\theta} \implies \\
&n/\overline{\theta} - (n-1)(1-\delta\gamma) \; > \; 1/\underline{\theta} \implies \\
&n\underline{\theta}/\overline{\theta} - (n-1)\underline{\theta}(1-\delta\gamma) \; > \; 1 \implies \\
&n\underline{\theta}(1 - \tfrac{(n-1)}{n}\overline{\theta}(1-\delta\gamma)) \; > \; \overline{\theta} \implies \\
&\underline{\theta}(1 - \tfrac{(n-1)}{n}\overline{\theta}(1-\delta\gamma))(1-\delta\gamma) \; > \; \bar{\theta}\frac{1-\delta\gamma}{n}
\end{aligned}
$$

Consequently, we have that $BR_i(\theta_i, e^\circ_{-i}, F_{-i}) > \bar{e}$ for every type θ_i. ∎

Lemma 5 *The extraction plan* $\mathbf{e}^\circ$ *that solves the indirect problem (46) is fully compressed.*

Proof of Lemma 5. Suppose, by contradiction, that e_i° is not compressed for some i. Then $\exists\ \theta_i',\ \theta_i''$ such that $\theta_i' < \theta_i''$ and $e_i^\circ(\theta_i') \neq e_i^\circ(\theta_i'')$. In fact, $e_i^\circ(\theta_i') < e_i^\circ(\theta_i'')$ by monotonicity. That is, incentive compatibility implies $\int_{\theta_{-i}} \log\left(\frac{1-\mathcal{E}^\circ(\theta)}{e_i^\circ(\theta)}\right) dF_{-i}$ is weakly decreasing in θ_i. So, symmetry of e° implies a point-by-point monotonicity restriction, namely $\log\left(\frac{1-\mathcal{E}^\circ(\theta)}{e_i^\circ(\theta)}\right)$ is weakly decreasing in θ_i. Consequently, $\frac{\partial e_i^\circ(\theta)}{\partial \theta_i} > 0$ holds for some i.

Without loss of generality, then, we take $\theta'' = \overline{\theta}$. From Lemma 3 and Lemma 4,

$$BR_i(\theta_i, e_{-i}^\circ, F_{-i}) > \bar{e} \ \geq \ e_i^\circ(\theta) \ \ \forall\ \theta.$$

Since i's preferences are single peaked in e_i, this means that

$$R_i^\circ(\overline{\theta}) - \theta_i'\ Q_i^\circ(\overline{\theta}) \ > \ R_i^\circ(\theta_i') - \theta_i'\ Q_i^\circ(\theta_i'),$$

for type $\theta_i' < \overline{\theta}$. In other words, i's interim expected payoff if it is type θ_i' is larger if it mimics $\overline{\theta}_i$. But this contradicts the incentive compatibility of $\mathbf{e}^\circ$. ∎

Lemma 6 *Consider the compressed extraction plan* $\mathbf{e}^\circ$ *where*

$$e_i^\circ = \frac{1}{n}(1-\delta\gamma)\int \theta_i dF_i \quad \forall\ i.$$

Then $\mathbf{e}^\circ$ *maximizes the indirect problem in (46) among all compressed extraction plans.*

Proof of Lemma 6. By Lemma 4, $Q_i^\circ(\theta_i) = Q_i^\circ(\underline{\theta})$ for all types θ_i. Now evaluate (45) at $\theta_i = \underline{\theta}$ and set $Q_i^\circ(\theta_i) = Q_i^\circ(\underline{\theta})$ for all θ_i as prescribed by full compression. This yields:

$$R_i^\circ(\underline{\theta}) - \overline{\theta} Q_i^\circ(\underline{\theta}) \ = \ R_i^\circ(\overline{\theta}) - \overline{\theta}\ Q_i^\circ(\overline{\theta})$$

Substituting this back into the indirect problem in (46) yields

$$\sum_i \left[R_i^\circ(\underline{\theta}) \ - \ Q_i^\circ(\underline{\theta}) \left(\overline{\theta} - \int_{\underline{\theta}}^{\overline{\theta}} F_i(\tilde{\theta}_i) d\tilde{\theta}_i \right) \right] \tag{53}$$

By a simple integration by parts, we obtain $\overline{\theta} - \int_{\underline{\theta}}^{\overline{\theta}} F_i(\tilde{\theta}_i) d\tilde{\theta}_i = \int_{\underline{\theta}}^{\overline{\theta}} \tilde{\theta}_i dF_i(\tilde{\theta}_i)$. We use this fact, together with the fact that $e_i(\theta)$ does not vary with type, to rewrite the objective in (46) as

$$\max \quad \sum_i \left[r(e^\circ(\underline{\theta})) \; - \; q_i(e^\circ(\underline{\theta})) \left(\int_{\underline{\theta}}^{\overline{\theta}} \tilde{\theta}_i dF_i(\tilde{\theta}_i) \right) \right] \tag{54}$$

Since the functions $r(e^\circ(\underline{\theta}))$ and $q_i(e^\circ(\underline{\theta}))$ are simply scalar choice variables of the planner, the solution to (54) yields the full information optimum in which each individual's type is $\int \theta_i dF_i$. The solution to this full information problem is the first best quota when i's type is $\int \theta_i dF_i$. Specifically, the solution is given by

$$e_i^\circ \;=\; \frac{(1-\delta\gamma)\left[\int_{\underline{\theta}}^{\overline{\theta}} \theta_i dF_i(\theta_i)\right]}{n}$$

This concludes the proof of Lemma 6. ∎

Combining Lemma 5 with Lemma 6, it follows that Part 1 of the Proposition holds. ∎

Proof of Part 2. Given the compressed quota system $\mathbf{c}^\circ$ from part 1, we construct a PBE (σ, μ) which implements $\mathbf{c}^\circ$. The proof largely mimics the steps of Proposition 1 which we do not repeat here. In particular, we construct a recursive sequence $\{e^\tau\}$ in the same manner, but now, each e^τ is constructed to be independent of θ. To do this, the incentive constraint for each τ requires that for all i,

(55)

$$\begin{aligned} V_i^{\tau-1} \;\geq\; \max_{\theta_i} \Bigg\{ & \frac{1-\delta\rho}{1-\delta\rho-\delta^2(1-\rho)} \;\max_{e_i} \left[\left(\frac{1}{1-\delta\gamma} - \theta_i \right) \log(1 - \mathcal{E}_{-i}^{\tau-1} + e_i) \;+\; \theta_i \log e_i \right] \\ & + \; \frac{\delta^2}{1-\delta\rho-\delta^2(1-\rho)} \left[\left(\frac{1}{1-\delta\gamma} - \theta_i \right) \log(1 - \mathcal{E}^\tau) \;+\; \theta_i \log e_i^\tau \right] \Bigg\} \end{aligned}$$

Hence, the difference between (55) and the analogous constraint (38) in the full information case is that the constraint here does not depend on the value of the realization of θ_i. As before, we can satisfy (55) by choosing $\mathcal{E}^\tau$ sufficiently close to one. Analogous to the full information case, the strategy profile (μ, σ) is constructed such that μ is truth-telling, and $\sigma_i(h^t(e^\tau), \tilde{\theta}, \theta_i) = \omega_t e_i^{\tau+1}$, $t \geq 1$ and $\sigma_i(h^0, , \tilde{\theta}, \theta_i) = \omega_0 e_i^\circ$, for all $\tilde{\theta}$ profiles disclosed, all types θ_i, and all countries i. By construction, the profile (μ, σ) is a Perfect Public Bayesian equilibrium that implements $\mathbf{e}^\circ$. ∎

7.4 Proof of Proposition 3

. First, recall the definitions of $r_i(\tilde{\theta})$ and $q_i(\tilde{\theta})$ and of R_i and Q_i given by equations (40), (41), and (42).

Now define the transfers $s_i(\theta)$ for all countries by

$$s_i(\tilde{\theta}_i, \theta_{-i}) \;=\; \int_{\theta_{-i}} \sum_{j\neq i} [r_j(\tilde{\theta}_i, \theta_{-i}) \\ -\; \theta_j q_j(\tilde{\theta}_i, \theta_{-i})]\; dF_{-i} \;-\; \frac{1}{n-1} \sum_{j\neq i} \int_{\theta_{-j}} \sum_{k\neq j} [r_k(\theta) \;-\; \theta_k q_k(\theta)]\; dF_{-j}$$

Observe that

$$\sum_i s_i(\tilde{\theta}_i, \theta_{-i}) \;=\; \sum_i \int_{\theta_{-i}} \sum_{j\neq i} [r_j(\tilde{\theta}_i, \theta_{-i}) \;-\; \theta_j q_j(\tilde{\theta}_i, \theta_{-i})]\; dF_{-i} \\ -\; \sum_i \frac{1}{n-1} \sum_{j\neq i} \int_{\theta_{-j}} \sum_{k\neq j} [r_k(\theta) \;-\; \theta_k q_k(\theta)]\; dF_{-j} \;=\; 0$$

so that these transfers balance ex post. To verify IC, we need to show

(56)

$$R_i(\theta_i) - \theta_i\, Q_i(\theta_i) + \int_{\theta_{-i}} s_i(\theta) dF_{-i} \;\geq\; R_i(\tilde{\theta}_i) - \theta_i\, Q_i(\tilde{\theta}_i) + \int_{\theta_{-i}} s_i(\tilde{\theta}_i, \theta_{-i}) dF_{-i} \quad \forall\, \theta_i \;\forall\, \tilde{\theta}_i$$

Observe that if i report type $\tilde{\theta}_i$, then

$$R_i(\tilde{\theta}_i, \theta_{-i}) \;-\; \theta_i\; Q_i(\tilde{\theta}_i, \theta_{-i}) \;+\; \int_{\theta_{-i}} s_i(\tilde{\theta}_i, \theta_{-i}) dF_{-i}$$

$$= \int_{\theta_{-i}} [r_i(\tilde{\theta}_i, \theta_{-i}) \;-\; \theta_i\; q_i(\tilde{\theta}_i, \theta_{-i}) \;+\; s_i(\tilde{\theta}_i, \theta_{-i})] dF_{-i}$$

$$= \int_{\theta_{-i}} \left\{ r_i(\tilde{\theta}_i, \theta_{-i}) \;-\; \theta_i\; q_i(\tilde{\theta}_i, \theta_{-i}) + \sum_{j\neq i} [r_j(\tilde{\theta}_i, \theta_{-i}) \;-\; \theta_j q_j \tilde{\theta}_i, \theta_{-i})] \right\} dF_{-i} \\ -\; \frac{1}{n-1} \sum_{j\neq i} \int_{\theta_{-j}} \sum_{k\neq j} [r_k(\theta) \;-\; \theta_k q_k(\theta)]\; dF_{-j}$$

Notice that the second term of s_i (the term beginning with $\frac{1}{n-1}$) is independent of i's reported type $\tilde{\theta}$, and so we can drop it from the notation. Hence, i's disclosure-relevant payoff is

$$\int_{\theta_{-i}} \sum_j [r_j(\tilde{\theta}_i, \theta_{-i}) \;-\; \theta_j q_j(\tilde{\theta}_i, \theta_{-i})] dF_{-i}$$

But this is simply the planner's payoff for a planner that knows only i's type. Yet, recall, that the full information quota choice for i does not depend on the realized types of others. Hence, if r and q correspond to the full information optimal quota system e°, then the payoff above is maximized by setting $\tilde{\theta}_i = \theta_i$. This would yield $e_i(\theta_i) = \theta_i(1 - \delta\gamma)/n$ which is the socially optimal choice for a type θ_i regardless of the distribution over θ_{-i}. ∎

References

[1] Acemoglu, D., P. Aghion, L. Bursztyn, and D. Hemous (2012), "The Environment and Directed Technical Change," *American Economic Review*, 102: 131-166.

[2] Amador, M., G.-M. Angeletos, and I. Werning (2006): "Commitment Vs. Flexibility," *Econometrica*, 74: 365-396.

[3] Antoniadou, E., C. Koulovatianos, and L. Mirman (2013), "Strategic Exploitation of a Common Property Resource under Uncertainty, *Journal of Environmental Economics and Management*, 65: 28-39.

[4] Aoyagi, M. (2003), "Bid Rotation and Collusion in Repeated Auctions," *Journal of Economic Theory*, 112: 79-105.

[5] Athey, S., A. Atkeson, and P. Kehoe (2005), "The Optimal Degree of Discretion in Monetary Policy," *Econometrica*, 73: 1431-1475.

[6] Athey, S. and K. Bagwell (2008), "Collusion with Persistent Cost Shocks," *Econometrica*, 76: 493-540.

[7] Athey, S., K. Bagwell, and C. Sanchirico (2001), "Collusion and Price Rigidity," *The Review of Economic Studies*, 17: 317-49.

[8] Arava, R., Y. Narahari, Deepak Bagchi, P. Suresh, S. V. Subrahmanya (2010), "Mechanism design problems in carbon economics," *Journal of the Indian Institute of Science*, 90:3 Jul-Sep.

[9] Baliga, S. and E. Maskin (2003), "Mechanism Design for the Environment," in *Handbook of Environmental Economics*, Volume 1, K.-G. Muller and Z. R. Vincent, Eds., Elsevier Science B.V. A.

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

[11] Battaglini, M. and Bard Harstad (2012), "Participation and Duration of Environmental Agreements," NBER working paper, #18585.

[12] Bodansky, D. (2004), "International Climate Efforts Beyond 2012: A Survey of Approaches," *PEW Center on Global Climate Change.*

[13] Cai, Y., K. Judd, and T. Lontzek (2012), "Tipping Points in a Dynamic Stochastic IAM," mimeo, January.

[14] Burke, M., J Dykema, D. Lobell, E. Miguel, and S Satyanath (2011), "Incorporating Climate Uncertainty into Estimates of Climate Change Impacts, with Applications to U.S. and African Agriculture," NBER Working paper #17092.

[15] Calcott, P. and V. Petkov (2012), "International Environmental Agreements: Design of Optimal Transfers Under Heterogeneity," *Environmental Modeling and Assessment*, 17: 209-220.

[16] Cave, J. (1987), "Long-Term Competition in a Dynamic Game: The Cold Fish War," *Rand Journal*, 18: 596-610.

[17] d'Aspremont, C., and L.A. Gerard-Varet (1979), "Incentives and incomplete information," *Journal of Public Economics*, 11: 25- 45.

[18] Desmet, K. and E. Rossi-Hansberg (2013), "On the Spatial Economic Impact of Global Warming," mimeo, Princeton University, February.

[19] Doepke, M. and R. M. Townsend (2006), "Dynamic Mechanism Design with Hidden Income and Hidden Actions," *Journal of Economic theory*, 126: 235-85.

[20] Dutta, P. and R. Radner (2009), "A Strategic Analysis of Global Warming: theory and some numbers," *Journal of Economic Behavior and Organization*, 71: 187-2009.

[21] Dutta, P. and R. Radner (2006), "A Game-Theoretic Approach to Global Warming," *Advances in Mathematical Economics*, 8: 135-53.

[22] Dutta, P. and R. Radner (2004), "Self-enforcing climate-change treaties," *Proceedings of the National Academy of Sciences.*

[23] Dutta, P. and R. K Sundaram (1998), "The Equilibrium Existence Problem in General Markovian Games," in *Organizations with Incomplete Information: Essays in Economic Analysis, A Tribute to Roy Radner*, M. Majumdar, Ed., Cambridge, UK: Cambridge University Press.

[24] Finus, M. (2001), *Game Theory and International Environmental Cooperation*, Cheltenham, Edward Elgar.

[25] Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski (2014), "Optimal Taxes on Fossil Fuel in General Equilibrium," *Econometrica*, 82: 41-88.

[26] Halac, M. and P. Yared (2014), "Fiscal Rules and Discretion under Persistent Shocks," *Econometrica*, 82: 1557-1614.

[27] Harrison, R. and R. Lagunoff (2016), "Tipping Points and Business-as-Usual in the Global Carbon Commons," revised, January 2016.

[28] Haurie, A. (2008), "A Dynamic Game Model of Strategic R&D Cooperation and GHG Emission Mitigation," mimeo.

[29] Krusell, P. and A. Smith (2009), "Macroeconomics and Global Climate Change: Transition for a Many-Region Economy," working paper.

[30] Levhari and L. Mirman (1980), "The Great Fish War: An Example Using a Dynamic Cournot-Nash Solution," *Bell Journal of Economics*, 14: 322-34.

[31] Lewis, T. (1996), "Protecting the Environment when Costs and Benefits are Privately Known," *Rand Journal of Economics*, 27: 819-47.

[32] Long, N.V. (2011), "Dynamic Games in the Economics of Natural Resources: A Survey," *Dynamic Games and Applications* , 1: 115-148.

[33] McAfee, R.P. and J. McMillan, "Bidding Rings," *American Economic Review*, 82: 579-99.

[34] Myerson, R. (1981), "Optimal Auction Design," *Mathematics of Operations Research*, 6: 58-73.

[35] Nordhaus, W. (2006), "After Kyoto: Alternative Mechanisms to Control Global Warming," *American Economic Review*, 96: 31-34.

[36] Nordhaus, W. (2007), "To Tax or Not to Tax: The Case for a Carbon Tax," *Review of Environmental Economics and Policy.*

[37] Ostrom, E. (1990), *Governing the Commons: The Evolution of Institutions for Collective Action,* Cambridge University Press, New York.

[38] Ostrom, E. (2002), "Common Pool Resource and Institutions: Towards a Revised Theory," *Handbook of Agricultural Economics, Volume 2.*

[39] Pavan, A., I. Segal, and J. Toikka (2014), "Dynamic Mechanism Design: a Myersonian Approach," *Econometrica.*

Dynamic Mech Design:

[40] Rouillon, S. (2010), "Optimal decentralized management of a natural resource," mimeo, Universite Bordeaux.

[41] Roberts, K. (1984), "The Theoretical Limits of Redistribution," *Review of Economic Studies*, 51(2): 177-95.

[42] Skrzypacz A. and H Hopenhayn (2004), "Tacit Collusion in Repeated Auctions," *Journal of Economic Theory*, 114: 153-69.

[43] Stern, N. (2006), *The Stern Review Report on the Economics of Climate Change.*

[44] United Nations Inter-Governmental Panel on Climate Change (2007), *Fourth Assessment Report: Climate Change.*