Dynamic Mechanism Design for a Global Commons

Rodrigo Harrison* and **Roger Lagunoff†**

Preliminary Draft

April 16, 2013‡

Abstract

This paper studies dynamic mechanisms for a global commons with environmental externalities. A leading example is carbon consumption. At each date, each country benefits from both from the use and the aggregate conservation of an open access resource. Conservation is beneficial because it reduces a country's environmental costs from consumption. The relative benefits of consumption compared to conservation are summarized by a privately observed parameter — the country's resource type — which evolves stochastically each period. An *optimal quota system* is an international agreement over resource consumption that maximizes world welfare subject to the constraint that it be implementable by Perfect Bayesian equilibrium compliance and disclosure strategies. Not surprisingly, with complete information the optimal quota allocates more of the resource each period to those countries with high value of consumption (and low value for conservation). However, under incomplete information, we show that the optimal quota is invariant to the country's resource type at every point in time. In the case of CO_2, this means that all *ex ante* identical countries receive the *same* emissions restrictions despite having differently evolved environmental costs and resource needs. We refer to this property as *extreme quota compression*, and show that extreme compression is robust to the distributional process on private shocks.

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

Key Words and Phrases: Dynamic mechanism design, global commons, optimal quota system, 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.

‡Initial draft.

1 Introduction

This paper examines the problem of mechanism design in a global commons framework. We model of a collection of countries that regularly consume a global, open-access resource. The resource is depletable, and its aggegate use imposes environmental costs on each country.

Examples include ozone depletion, ocean fisheries, and deforestation. A leading example is atmospheric emission of CO_2. The use of carbon based resources results in increased concentrations of green house gases (GHG) that impose costs on a country's economy through its effects on climate. Estimates provided to the United Nations Inter-governmental Panel on Climate Change (IPCC) range from 1 to 5% in global GDP reduction due to an increase of 4°C.[1]

There is, by now, a large literature analyzing mechanisms addressing global commons problems such as GHG emissions. Much of it focusses on a fairly narrow range of practical options, including variations of cap and trade, carbon taxes, carbon credit exchanges, and other well publicized proposals. With some exceptions, the central concern motivating the discussion of these options is productive efficiency — finding the most cost-effective market mechanism to allocate carbon. Informational incentives, the classic concern of formal models of economic mechanisms, have received less attention.[2]

Despite some early contributions such as Nordhaus and Yang (1996), Baliga and Maskin (2003), and Dutta and Radner (2006, 2009), there are strikingly few formal models of unencumbered mechanism design for the global commons.[3]

One reason for this is that many global commons environments present an unusual combination of challenges. Using carbon consumption as an example, the global scale of GHG emissions means that many if not most countries must be involved in the negotiations. Since there is obviously no international government to enforce limits, any viable mechanism must be *dynamically* self-enforcing in both compliance and in truthful revelation of information. To complicate matters further, the accumulation of atmospheric CO_2 is an inherently dynamic process that creates an accretive, negative externality (climate change). Its effects are difficult to predict and are heterogeneous across countries. According to the 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

[1] IPCC Fourth Assessment Report: Climate Change 2007.

[2] Arava, et. al. (2010) provide an excellent summary of the mechanisms in place and their rationales.

[3] See Section 2 for a fuller discussion of the literature.

> populations and very likely underestimate damage costs because they cannot include many non-quantifiable impacts. ..” . (IPCC 2007 Synthesis Report).

Given all this, any international agreement must be structured so that countries find it in their self interest to follow its prescriptions, all the while accounting for difficult-to-predict changes in the benfits and costs of carbon usage, and in any other new, possibly asymmetric, information as it arrives.

To address the issues in a tractable way, the present paper posits an infinite horizon model of international resource consumption. 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 in the model because it allows the country to avoid the environmental costs of aggregate resource extraction and consumption. The conservation benefits are heterogeneous across counties and are assumed to evolve stochastically as countries are hit with private, idiosyncratic “payoff” shocks each period. These shocks may be serially correlated, and the distributions across countries may differ due to geographic and demographic influences. The shock process captures a common feature of many commons problems: environmental costs are often difficult to forecast and typically vary widely across countries.

The model generalizes the well known “Fish War” games originating with Levhari-Mirman (1980).[4] In their classic model, Levhari and Mirman study the strategic allocation of a depletable, open access resource such as fish or forests. Identical users in their model choose how much to consume each period, leaving the residual for future extraction. There are no direct costs or externalities from usage. Conservation is therefore valued in the fish war game only for instrumental reasons: preserving the stock allows one to smooth consumption.

This paper modifies the Fish War framework by adding a heterogeneous “usage externality” that makes conservation beneficial. We refer to the parameter that determines the country’s value of conservation (relative to its use) as its *resource type*. We then build in serially correlated private shocks to each country’s resource type each period.

A *quota system* is an international agreement that specifies limits for each country’s resource consumption (or emissions) at each point in time, given the current stock and given the payoff characteristics of all countries. For a quota system to be feasible, however, it must be implementable by a Perfect Bayesian equilibrium (PBE). In a PBE, each country optimally chooses its disclosure and its level of resource consumption each period, given its updated beliefs after observing the public disclosure and usage history available at the time. Implementability in PBE is a natural requirement as it takes account of the strategic incentives of countries to follow through on the agreement, given their private information.

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

We focus on the characteristics of an *optimal quota system*, that is, a quota system that jointly maximizes the expected long run payoffs of all countries such that it can be implemented by a PBE.

The benchmark case is one with full information. Each country's payoff for consumption/ conservation is known and there are no shocks. The optimal quota in this case is easily characterized by stationary usage rates that vary across countries. Those countries that place high value on consumption (or low value on conservation) are permitted to extract more.

Interestingly, the optimal quota is not necessarily implementable by the threat of a "fish war," even at arbitrarily high discount factors. The "fish war" in this case refers to the heterogeneous extension of Levhari-Mirman's result, namely, the unique Markov Perfect equilibrium in which each country over-extracts relative to the optimum. We show for certain conditions on type profiles, that the threat of reversion to a Fish War *does not work to implement the optimal quota*, even for discount factors close to one. Certain extreme types may prefer the fish war to the social optimum. Nevertheless, we show that the optimal quota *can be* implemented by a PBE that uses graduated punishments that further deplete the stock each time a country violates its prescribed resource use.

The main results pertain to the case of incomplete information — the case where persistent, private payoffs shocks hit each of the countries each period. Under private information, all countries have incentives to chooses extraction policies that overstate their values for extraction. Hence, under private information, the constrained-optimal quota system is not efficient, *ex post*.

We first analyze a special case where the shocks are perfectly persistent; each country privately draws its realized payoff type once and for all at date $t = 0$. We show that the optimal quota in this case is completely insensitive to a country's realized type. Unlike in the full information optimum, the optimal quota is not tailored to the realized benefits and costs of resource usage. We refer to this as the property of *full* or *extreme quota compression*.

To illustrate what this means, consider two countries that draw their types from the same distribution, ex ante. Extreme compression implies that the quota for, say, carbon allocates too little of the resource to the country that experiences higher than average carbon (fossil fuel) needs and/or lower-than-average costs to climate change. By contrast the quota allocates an excessively high limit to the country with lower than average resource needs and/or higher-than-average costs to climate change.

The result is reminiscent of key research by Athey and Bagwell (2008) who study optimal collusion in oligopolies in which firms incur serially correlated cost shocks.[5] They show that under a monotone hazard rate restriction on the distribution of shocks, the optimal allocation of each firm's market share is independent of its realized cost type (a property which they

[5]See also Athey, Bagwell, on Sanchirico (2001), Aoyagi (2003), and Skyrpazc and Hopenhayn (2004) for related models with iid shocks.

refer to as "rigidity").

Our notion of quota compression clearly resembles their notion of market share rigidity. It turns out, however, compression holds for somewhat different reasons and in different circumstances. Here, extreme compression holds regardless of the shape of the shape of the hazard rates and, more generally, for fairly arbitrary distributions on shocks.

To understand why it holds in the global commons problem, the comparison with a collusive oligopoly or a procurement auction is instructive. As in these environments, every mechanism (quota system) in the global commons environment also has a "market share" effect and a "compensation" effect on each participant's long run payoff. By standard envelope arguments, the equilibrium payoff is shown to depend only on the "market share" effect. When the participants are firms, this means that optimal mechanisms will not generally be compressed because the allocation of market share inherently requires trade-offs, either between the firms and the buyer-designer in the procurement auction, or among the firms themselves in the collusive oligopoly. In the latter environment, a higher production quota assigned to one type of seller must be offset with a lower one to another. The oligopoly would ideally like to increase all firms' quotas but cannot.

By contrast, in the global commons problem , a country's "market share" is its expected net present value of "stored resource". A country's expected net present value of "stored carbon", for instance, is the present value it places on all carbon currently stored in fossil fuels and in forest cover, relative to its carbon consumption. Because stored carbon is like a public good, if the planner wishes to increase each country's value of stored carbon at once, it can do so merely by reducing everyone's usage quota. Thus there is no allocative trade off between participants' "market share" of stored carbon.

As this intuition suggests, the result depends critically on there being a public good-like attribute to the commons problem. If, say, the countries derived no direct benefit from conservation and the resource was a purely private good (as in the classic fish war model), then private information would not necessarily lead to quota compression.

We characterize the compressed quota, and further show that the compressed quota can be implemented by a PBE which itself uses graduated punishments that are also compressed. The compression of punishments, it turns out, is needed to prevent double deviations in both disclosure and in subsequent resource use.

We are able to extend this result to the case of imperfectly persistent shocks. Once again, the quota exhibits extreme compression. In this case, however, the quota need not be time stationary. Because the distributions can allow for drift, the quota can become more severe over time for all countries if, for instance, the costs of climate change are expected to increase. Hence, the optimal quota can exhibit time variation but not cross-country variation.

The rest of the paper is organized as follows. The next section (Section 2) summarizes the mechanism design literature as it applies to carbon limits and resource extraction more generally. We then relate the present model, to the recent theoretical innovations in dynamic mechanism design. 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. We first take up the case of perfect persistence, and then extend the results to the general, imperfect persistence case. Section 6 concludes with a discussion of how quota compression fits in to the policy framework proposed under the UN Framework Convention on Climate Change.

2 Literature on Dynamic Climate Mechanisms

A common type of analysis of the global commons is one that analyzes "pros" and "cons" of many of the practical options discussed in policy circles. In the climate discussion, in particular, a central concern is on finding the most cost-effective way to acheive a fixed carbon target.

While the quantitative and policy literature is far too large to survey here, we refer the reader to the references contained in the following recent contributions, all of which contain more comprehensive surveys. Arava, et. al. (2010) offers a useful description and classification of the various proposals ratified under Kyoto. Some key quantitative assessments of particular mechanisms include Nordhaus (2006, 2007) and Stern (2006) and Golosov et al (2011), all of whom offer quantitative assessments of carbon tax policies. 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. Bodansky (2004) discusses the wide variety of factors and constraints that should appear in a full elaboration of the global mechanism design problem for carbon allocation.

Barrett (2003) and Finus (2001) argue that any international mechanisms must be self-enforcing. This lead them to propose repeated game models in which international climate agreements are implementable in subgame perfect equilibria.

A number of papers have proposed models that extend the self-enforcement constraint to non stationary commons games that better characterize the dynamics of resource use.[6] Rouillon (2010) examines the classic fish war model of Lehvari and Mirman (1980). He proposes a competitive pricing scheme with directed transfers between the players that operates when the resource can be traded exchanged in a competitive market. Cave (1989) also examines the traditional full information fish war model. He showed that reversion strategies (i.e., using the threat of a Markov Perfect "Fish War" as punishment) can enforce full cooperation of the agreement when the participants are sufficiently patient.

[6] See also Ostrom (2002) for 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.

Dutta and Radner (2006, 2009) study a version of the fish war that adds a climate externality as we do. They (implicitly) characterize the optimal resource quota subject to the constraint that it be implemented by a Markov Perfect equilibria among the countries. Our study builds on the Dutta-Radner approach by studying informational incentives when countries receive persistent, heterogenous private shocks each period.

In a different type of dynamic game, Battaglini and Harstad (2012) examines endogenous coalition formation for solving 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 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.

In the aforementioned literature, there is no asymmetric information and so the self-enforcement constraint applies to questions of compliance and participation rather than of truthful disclosure. The present paper builds on these apprroaches by integrating both compliance and disclosure of private information into the dynamic constraints of any international agreement.

3 The Full Information Benchmark Model

3.1 Basic Setup

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

To better motivate the framework, we use carbon usage as a leading example. We acknowledge at the outset that the model does not fit the technology for carbon usage perfectly. In particular, the leading source of GHG emissions is fossil fuel which is not a purely open access resource, strictly speaking. For our purposes, however, the model is a reasonable approximation since (1) access to all types of resources that produce GHG emissions are fairly widely dispersed among a large collection of countries; and (2) 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.

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 internal resource usage.[7] 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 $C_t = \sum_i c_{it}$. 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 resource depreciates exponentially at rate γ. However, $\gamma > 1$ allows for growth in the stock, as in the case of carbon sequestration.[8] The transversality condition $\delta\gamma < 1$ is assumed to hold.

Let $\mathbf{c}_t = (c_{1t} \ldots, c_{nt})$ denote the t period profile of resource consumption levels. The entire dynamic path of resource consumption for all countries is the given by

$$\mathbf{c} \;=\; \{\mathbf{c}_t\}_{t=0}^{\infty}$$

Given a consumption path $\mathbf{c}$, the long run payoff to country i is given by

$$U_i(\omega_0, \mathbf{c}, \theta_i) \;\equiv\; \sum_{t=0}^{\infty} \delta^t \left[\theta_i \log c_{it} + (1-\theta_i) \log(\omega_t - C_t)\right] \tag{2}$$

Two Interpretations. Using carbon as the example, a country's payoff U_i in (2) can be interpreted in one of two ways. This first is to associate U_i simply with the preferences of a "representative citizen." The citizen's flow payoffs are discounted by δ each period. His flow payoff weights both resource consumption and resource conservation. The value θ_i is a pure preference weight given to one's own log consumption, whereas $1 - \theta_i$ is the weight assigned to the remaining stock $\omega_t - C_t$ of, for instance, usable carbon. The specification is motivated by the idea that 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 parameter θ_i will be referred to as i's "resource type" or simply its "type" and is assumed to lie in an interval $[\underline{\theta}, \overline{\theta}] \subset [0, 1]$. 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 varying types θ_{it} for each country i.

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

[7] The stock ω_t can also be interpreted to be the amount that, if fully depleted, would lead to to a loss of environmental sustainability.

[8] Sequestration refers to the natural process by which carbon is broken down by plants or re-absorbed into the oceans. It returns in either case to its "stored" state.

models assume $\theta_i = 1$ for all i, in which case a user of the resource merely trades off the value of immediate usage against the value of 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.[9] 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. which differs in its effect 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 intensity 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 representative consumer from country i at date t. The composite good is produced using two inputs, one carbon-based and the other not, according to the technology, $y_{it} = c_{it}^{\theta_i}(\omega_t - C_t)^{1-\theta_i}$.

According to this formulation, both inputs are produced from a basic resource. The carbon input gets used up in the production process, while the non-carbon input is renewable, but depreciates (or 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-based resource to produce a given output. Richer countries, for instance, have larger carbon requirements as a consequence of a more developed economy.

In either interpretation, 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.[10]

Finally, note that the model "abstracts away" issues of endogenous technical change and technology transfer between countries. Though these are clearly central issues in current discussions of climate mechanisms, the present study focuses at this stage purely on issues of disclosure and compliance.

[9] 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.

[10] Size differentials would be captured instead by differential welfare weighting in any planner's problem. We return to this issue when the planner's problem is introduced.

3.2 Quota Systems

A *global type profile* is given by $\theta = (\theta_1, \ldots, \theta_n) \in [\underline{\theta}, \overline{\theta}]^n$. Following standard notational convention, $\theta_{-i} = (\theta_j)_{j \neq i}$.

Our interest is in understanding the nature of international agreements that could recommended by an international agency given the type profile of its member countries. The international agency (IA) as envisioned here is a coordinating institution such as the U.N. It operates by the consent of its members, gathers and makes available information, makes recommendations, and suggests sanctions for violations (though it does not have the ability to enforce sanctions).

The object of choice for the IA is a *quota system*. A quota system is a dynamic path of resource consumption. It is defined formally as a consumption path $\mathbf{c}^*(\theta)$ that associates a type profile to a recommended consumption path for each country so that $c^*_{it}(\theta)$ is the consumption recommended for country i at date t given the global profile θ. For a quota system $\mathbf{c}^*(\theta)$ to be feasible, it must first be consistent with the intertemporal resource constraint $\sum_i c^*_{it}(\theta) \leq (\omega_t - \sum_i c^*_{i\ t-1}(\theta))^\gamma$ for all t. Second, the quota must be *implementable* by an equilibrium of the dynamic resource game.

To be clear about this second constraint, let $h^t = (\omega_0, \mathbf{c}_0, \omega_1, \mathbf{c}_1, \ldots, \omega_{t-1}, \mathbf{c}_{t-1}, \omega_t)$ denote the date t history of usage levels and resource stocks, including the current stock ω_t. The initial history is $h^0 = \omega_0$. Let H denote the set of all histories over all dates t. Then, a *usage strategy* $\sigma_i(h^t, \theta) = c_{it}$ for country i maps histories and global type profiles to resource consumption at date t. A *usage profile* is given by $\sigma = (\sigma_1, \ldots, \sigma_n)$.

Any given strategy profile σ may be said to *implement* a quota system $\mathbf{c}^*$ as follows. Starting at the initial date, let $\mathbf{c}^*_0(\theta) = \sigma(h^0, \theta)$, then define $\mathbf{c}^*_1(\theta) = \sigma(h^0, \mathbf{c}^*_0(\theta), (\omega_0 - C^*_0(\theta))^\gamma,\ \theta)$, and so forth... so that $\mathbf{c}^*$ merely describes the on-path consumption realized by σ.

3.3 Optimal Quota Systems

Since the IA is cannot directly impose or enforce a quota system, the system must be implemented by a subgame perfect equilibrium (SPE) profile of usage strategies.

A given profile σ is an SPE if for each country i, σ_i maximizes country i's long run payoff after each history h^t, given the type profile θ and the strategies σ_{-i} of others (see the Appendix for the formal recursive expression of continuation payoffs).

In its role as coordinating institution, the IA therefore recommends a quota $\mathbf{c}^*(\theta)$ and a subgame perfect profile σ that solves

$$\max_{\mathbf{c}^*} \sum_{i=1}^{n} U_i(\omega_0, \mathbf{c}^*(\theta), \theta) \quad \text{such that } \mathbf{c}^* \text{ is implemented by the SPE } \sigma. \tag{3}$$

Hence, the IA chooses the quota to maximize the joint sum of all countries' payoffs such that the quota be sequentially and credibly self-enforcing. The formulation in (3) 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 for our purposes this would complicate the notation without adding to the results.

The solution to (3) can be found by breaking the problem into two steps. Step 1 characterizes the optimal quota without the equilibrium constraint, i.e., i.e., without the requirement that the quota be implemented by a subgame perfect equilibrium. This "relaxed" problem would suffice if the IA could impose and enforce a quota system on its members. Step 2 shows that the "relaxed" solution can be implemented by SPE profile σ^*. This simple two-step algorithm will be repeated when private shocks are later introduced into the model.

Step 1 characterizes the solution to the relaxed problem — the optimal quota in the absence of the equilibrium constraint. It turns out to be easier to work with extraction *rates* rather than levels. For any dynamic path $\mathbf{c}$ of resource consumption, let $\mathbf{e}$ denote the corresponding extraction rates, as defined $e_{it} = \frac{c_{it}}{\omega_t}$. Let $\mathcal{E}_t = \sum_i e_{it}$ denote the aggregate extraction rate. Using rates rather than levels in the recursive payoffs for countries, the relaxed solution may be found using standard techniques to solve the IA's Euler equation derived from the sum of each country's payoff (each given by (2)),

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

(Note that one could redefine $\mathbf{e}$ so that it explicitly depended on the state, but since the effects on future rates are eliminated by Envelope arguments, we omit the notation for brevity.) Differentiating the value function for a country i gives

$$\frac{\partial U_j(\omega_t, \mathbf{e}, \theta)}{\partial \omega_t} \; = \; \frac{1}{\omega_t} + \delta\gamma \frac{\partial U_j}{\partial \omega_{t+1}} \omega^{\gamma-1} (1-\mathcal{E}_t)^{\gamma}$$

Iterating this second equation forward one period, then substituting it into the Euler equation on the right-hand side of the (iterated) expression gives, after some manipulation

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

Solving the forward equation yields

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

which is obviously stationary. Re-arranging terms and aggregating over i yields the aggregate rate $\mathcal{E}^*(\theta) = \frac{\Theta(1-\delta\gamma)}{n}$ where $\Theta \equiv \sum_i \theta_i$. In other words, the optimal aggregate usage is a fraction $(1-\delta\gamma)$ of the average resource type $\frac{\Theta}{n}$. This aggregate rate is achieved by country-specific rates $e_i^*(\theta) = \frac{\theta_i(1-\delta\gamma)}{n}$. These rates yields a quota system $\mathbf{c}^*$ given by

$$c_{it}^*(\theta) = \frac{\theta_i(1-\delta\gamma)}{n}\, \omega_0^{\gamma^t} \left(1 - \frac{\Theta(1-\delta\gamma)}{n}\right)^{\gamma(1-\gamma^t)/(1-\gamma)} \tag{4}$$

for country i in date t.

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 are increasing in one's own resource type θ_i but decreasing in the cross-country average $\frac{\Theta}{n}$. Hence, countries with relatively larger consumption value should be allowed to extract more. In other words, "pro-consumption" types should extract more while "pro-conservation" types should extract less.

The quota system $\mathbf{c}^*$ is then the optimal one if it can be implemented by a SPE. This is Step 2 in the characterization of (3) which we summarize in the following lemma.

Lemma 1 (Full Information Benchmark) *For any profile θ of resource types, the IA's optimal quota system (the solution to (3)) is the quota system $\mathbf{c}^*$ described by (4).*

The proof is given in the Appendix. Lemma 1 will serve as a useful benchmark against which later results may be compared. The result may be of independent interest because it holds for *any* discount factor $\delta > 0$, including impatient ones.[11] This comes from the fact that payoffs are not bounded below. The unbounded payoffs reflect the idea that in the case of global commons, the costs of resource depletion may be catastrophic.

A consequence of the unbounded payoffs is that 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 depends on even harsher punishment if the countries fail to carry out the sanction. This means, in turn, increasingly severe resource depletion must be threatened, itself made credible only by the threat of still further depletion later on, and so forth. The

[11] We are aware of one other result that sustains the optimal extraction rates in the common pool resource environment independently of discount factor. Rouillon (2010) examines the traditional fish war environment (i.e., no heterogeneity) and proposes an interesting pricing scheme with directed transfers between the players. The optimal extraction policy is sustained for any discount factor δ. Certain requirements in his setup limit its applicability to global commons problems between countries, however, since users are assumed to be price takers in a competitive market. In addition, it is unclear how the sanctions might work off-path in the competitive market model.

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.

Another virtue is that the equilibrium construction does not actually require that the IA monitor the individual resource usage of each country, 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.

3.4 Optimal Quota versus Fish War

The downside of the equilibrium in Lemma 1 is that it requires substantial intervention from a coordinating body. At each (counterfactual) stage, the IA needs to monitor usage, carry out the randomization, and determine the punishment at the next stage. All this coordination is unnecessary if the quota is enforced by the threat of a "fish war". The fish war, characterized by the unique, Markov Perfect equilibrium (MPE) of the game, requires no special coordination, no monitoring beyond the initial quota, and no especially harsh sanctions since MPE is precisely what would occur in a purely decentralized, "business-as-usual" setting. Under the MPE, countries do not cooperate in any attempt to jointly condition their usage strategies on past outcomes.

Though the MPE has no special significance in our analysis, it has been the subject of much attention in the resource literature, and given its advantages we study its viability as a credible punishment threat if the quota is violated.

Using the same standard Euler equation techniques as before, the Markov Perfect equilibrium (MPE) may be shown to implement a quota system that is stationary in usage rates for each country i, with rates given by $e_i^{\circ}(\theta) = \frac{\theta_i(1-\delta\gamma)}{1+\Theta_{-i}(1-\delta\gamma)}$ where $\Theta_{-i} = \Theta - \theta_i$. The MPE consumption levels are linear in the stock since $\sigma_i(\omega_t, \theta) = \omega_t e_i^{\circ}$. The MPE implements a quota system denoted here by $\mathbf{c}^{\circ}$ such that

$$c_{it}^{\circ}(\theta) \;=\; \frac{\theta_i(1-\delta\gamma)}{1+\Theta_{-i}(1-\delta\gamma)}\;\omega_0^{\gamma^t}\left(1-\sum_j \frac{\theta_j(1-\delta\gamma)}{1+\Theta_{-j}(1-\delta\gamma)}\right)^{\gamma(1-\gamma^t)/(1-\gamma)} \tag{5}$$

As with the optimal quota, the MPE accounts for heterogeneity by allowing pro-extraction types to extract more. It is easy to verify that usage rates are higher in the MPE, i.e., $e_i^*(\theta) > e_i^{\circ}(\theta)$. Hence, countries over-extract the resource in the fish war.

Unlike in the standard homogeneous fish war, however, it is not necessarily the case that all

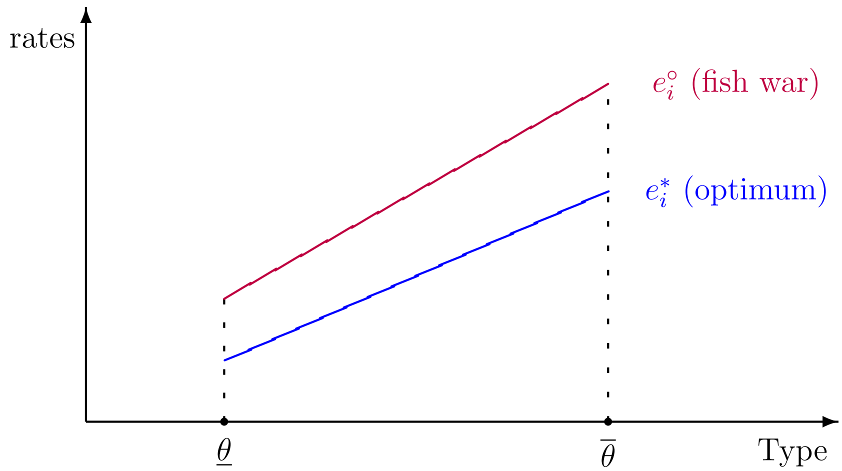

Figure 1: Comparison of Optimal Quota and the Fish War

countries are better off in the optimal quota. To see this, one can verify that a necessary condition for equilibrium supported by the threat of a fish war is: $V_i(\omega_t, \mathbf{c}^*(\theta), \theta) > V_i(\omega_t, \mathbf{c}^\circ(\theta), \theta)$, i.e., long run payoffs are higher under the optimal quota for all players than in the fish war. In fact, this only holds for certain parameters and in the limit as δ goes to one.

Lemma 2 (Implementation by Fish War) *Supposes the resource profile θ satisfies*

$$\frac{\Theta}{n} \geq \theta_i \frac{\log n}{n-1} \tag{6}$$

Then there exists a discount factor $\underline{\delta}$ and resource intensity $\underline{\gamma}$ such that for all (δ, γ) pairs such that $\delta \in [\overline{\delta}, 1)$ and $\gamma \in [\underline{\gamma}, \frac{1}{\delta})$, the optimal quota $\mathbf{c}^*$ is implementable by the threat of a fish war, i.e., the threat of permanent reversion to the Markov Perfect equilibrium $\mathbf{c}^\circ$.

On the other hand, suppose that the inequality in (6) is violated for some profile θ. Then for this profile, there is no discount factor for which $\mathbf{c}^*(\theta)$ is implementable by the threat of a fish war.

If the inequality in (6) is violated, it means that there are profiles θ such that at least one country prefers the decentralized Markov equilibrium to the International Agency's optimal quota system. Such a case does not imply the optimal quota system cannot be implemented, as we already know from Lemma 1. It does, however, indicate that simple triggers are not be viable.[12] The result mirrors the analysis in Dutta and Radner (2009) in a somewhat different model. They compare the socially optimal emissions quotas for different welfare weights against the Markov Perfect equilibrium (which they refer to as "Business-as-usual"). They also find the MPE trigger can implement the optimum for some, but not all, welfare weights.

[12]This result might be relaxed if monetary transfers were used.

4 Persistent Private Shocks

In most cases, the full costs and benefits of resource usage are not known in advance, although some countries are more likely to have high costs than others. We consider the case in which countries incur privately observed stochastic shocks to their resource types.

The first part of this section (Sections 4.1 and 4.2) introduces the general model with private shocks. There, the international agency's optimal quota problem is stated subject to the constraint that it be implemented in Perfect Bayesian equilibria (PBE). The PBE constraint requires both self enforcing compliance and self enforcing disclosure. The second part (Section 4.3) takes up the simple case of (perfectly) persistent shocks — each country's initial shock at $t = 0$ remains in place for all time. The model with persistent shocks is a reasonable approximation to an environment where technological progress moves more slowly than environmental change. In the final part (Section 4.4), the perfect persistence assumption is then relaxed, and the case of imperfect serial correlation is analyzed.

4.1 The Setup

Each period, each country is hit with an idiosyncratic, privately observed shock to its resource type. These shocks could represent unforseen changes to the country's environmental costs, or alternatively to its relative resource intensity.

Let θ_{it} denote the realized resource type of country i. The distribution over θ_{it} is given by a Markov kernel, $F_i(\theta_{it}|\theta_{i\ t-1})$ that determines the current resource type given the country's resource type in the prior period. The distribution $F_i(\cdot|\theta_i)$ is assumed to have full support on $[\underline{\theta}, \overline{\theta}]$ and admits a conditional density $f_i(\cdot|\theta_i)$ for each $\theta_i \in [\underline{\theta}, \overline{\theta}]$.

The shocks are IID across countries, and while the Markov distribution F_i for country i is commonly known to all countries, each county's realized shock each period is privately observed. The informational asymmetry reflects the fact that each country typically has superior knowledge of own immediate resource needs and depletion costs.[13]

The profile of types in date t is $\theta_t = (\theta_{1t}, \theta_{2t}, \ldots, \theta_{it}, \ldots \theta_{nt})$. Let $F(\cdot|\theta_{t-1})$ denote the joint distribution on the type profile θ_t conditional on the type profile θ_{t-1}, and let $F_{-i}(\cdot|\theta_{-i\ t-1})$ the joint conditional distribution excluding country i.

We will find it useful to express the ex ante distribution of a shock realization t periods ahead. Let $\theta_i^t = (\theta_{i0}, \theta_{i1}, \ldots, \theta_{it})$ the history of realized types for country i. Then let $F_i^{t-\tau}(\theta_i^t|\theta_{i\tau})$ denote the $t - \tau$ period-ahead forecast of θ_i^t given $\theta_{i\tau}$. The initial prior is denoted

[13]International institutions tend, for the most part, to rely on most countries' national income accounting and information gathering to estimate economic costs and benefits of emissions restrictions.

by $F_i^0(\theta_{i0})$. Finally, let $\theta^t = (\theta_0, \theta_1, \ldots, \theta_t)$ be the history of profiles of all countries' realized types, and $F^t(\theta^t)$ the ex ante (date 0) distribution over history θ^t.

4.2 Optimal Quota Systems in the Private Shocks Model

As before, the international agency (IA) recommends a quota system. With private shocks, a quota system is now given by the sequence $\mathbf{c}^* = \{c_t^*(\theta^t)\}, t = 0, 1, \ldots$. Notice, in other words, that the IA's recommendation at each date now depends on the entire history of shocks θ^t up to that point. To see why this must be the case, observe that the initial recommendation c_0^* will depend on the initial realization θ_0. Next period's recommendation c_1^* will depend, of course, on θ_1. But in order for next period's recommendation to be feasible, the aggregate consumption C_1^* must be bounded by the state ω_1 which depends on last period's consumption which, in turn, depends on the initial shock θ_0. Rolling things back from an arbitrary date t, the current recommendation must depend on the realized path θ^t.

Letting $\omega_t^*(\theta^{t-1})$ express the resource stock on-path as a function of the past shock history, a country's expected payoff under quota system $\mathbf{c}^*$ can be expressed after any history θ^t:

$$U_i(\omega_t, c^*(\theta^t), \theta_{it}) \ = \ \sum_{\tau=t}^{\infty} \delta^{\tau-t} \int_{\theta^\tau} \left[\theta_{i\tau} \log c_{i\tau}^*(\theta^\tau) + (1 - \theta_{i\tau} \log(\omega_\tau^*(\theta^{\tau-1}) - C_\tau^*(\theta^\tau) \ ) \ dF^{\tau-t}(\theta^\tau|\theta_{it})\right] \tag{7}$$

where we use the convention $dF_i^t(\theta_{it}|\theta_{it}) = 1$.

To make an effective recommendation, the international agency's role is now expanded to include that of gatherer and dispenser of information. At the beginning of each period, the IA solicits information concerning each country's realized type θ_{it}. 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).[14]

As with resource consumption, the IA cannot enforce credible disclosure by its members. Nor can it commit in advance how the information will influence its recommendations. Instead, the IA serves as a vehicle for coordinating information and usage.

Denote each country's report by $\tilde{\theta}_{it}$. A entire profile of reports is denoted $\tilde{\theta}_t$. We will refer to $\tilde{\theta}_t = ((\tilde{\theta}_0, \tilde{\theta}_1, \ldots, \tilde{\theta}_t)$ as the *disclosure history.* Finally, let $\theta^t \backslash \tilde{\theta}_{i\tau}$ denote a t disclosure history in which with $\tilde{\theta}_{i\tau}$ is substituted for $\theta_{i\tau}$ at date $\tau \leq t$.

In order to obtain the desired quota system, the IA suggests a strategy profile that now includes a disclosure option. Formally a *disclosure strategy* for country i is map $\mu_i(h^t, \tilde{\theta}^{t-1}, \theta_{it}) =$

[14]This is also in line with existing international protocols. Article 12 of the UN Framework Convention for Climate Change, 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..."

$\tilde{\theta}_{it}$ is i's report to the IA in period t, given the usage history, the disclosure history, and i's current resource type.

The definition of usage strategy also needs slight modification to account for the disclosure history. Country i's usage strategy is a now map $\sigma_i(h^t, \tilde{\theta}^t, \theta_{it}) = c_{it}$ determining i consumption at date t given usage history, the disclosure history, and i's current resource type.

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 $\mu(h^t, \tilde{\theta}^{t-1}, \theta_t) = (\mu_i(h^t, \tilde{\theta}^{t-1}, \theta_{it}))_{i=1}^n$ and $\sigma(h^t, \tilde{\theta}^t, \theta_t) = (\sigma_i(h^t, \tilde{\theta}^t, \theta_{it}))_{i=1}^n$. Given a strategy pair (σ, μ), the long run expected payoff to a country i at the resource consumption stage in date t is

$$
\begin{aligned}
V_i(h^t, \tilde{\theta}^t, \sigma, \mu | \, \theta_{it}) \; \equiv \; \int_{\theta_{-it}} \Bigg[& \theta_{it} \log \sigma_i(h^t, \tilde{\theta}^t, \theta_{it}) + (1-\theta_{it}) \log(\omega_t - \sum_{j=1}^{n} \sigma_j(h^t, \; \tilde{\theta}^t, \theta_{jt}) \;) \\
& + \delta \int_{\theta_{i\,t+1}} V_i(h^{t+1}, \tilde{\theta}^{t+1}, \sigma, \mu | \; \theta_{i\,t+1}) dF_{-i}(\theta_{i\,t+1} | \theta_{it}) \Bigg] \; dF^*_{-i}(\theta_{-it} | \; h^t, \tilde{\theta}^t)
\end{aligned} \tag{8}
$$

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

At the disclosure stage, country i evaluates its payoff before observing the disclosed type of others. In this case its payoff is $\int_{\theta_{-it}} V_i(h^t, \tilde{\theta}^{t-1}, \mu(h^t, \tilde{\theta}^{t-1}, \theta_t), \sigma, \mu | \; \theta_{it}) dF^*_{-i}(\theta_{-it} | h^t, \tilde{\theta}^{t-1})$.

To implement a quota system, the IA suggests 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 $\tilde{\theta}_t$. We focus on truth-telling disclosure strategies, i.e., those in which μ prescribes $\tilde{\theta}_{it} = \theta_{it}$ for each country. The strategy pair (μ, σ) with truth-telling disclosure may then be said to *implement* the quota system $\mathbf{c}^*$ in the private shocks model if

$\mathbf{c}^*$ is induced by prescribed (on-path) play, that is,

$$\mathbf{c}_0^*(\theta^0) = \sigma(h^0, \theta^0,\ \theta_0),$$

$$\mathbf{c}_1^*(\theta^1) = \sigma(h^1, \theta^0, \theta_1), \quad \text{where } h^1 = (h^0, \sigma(h^0, \theta_0,\ \theta_0), (\omega_0 - C_0(\theta^0))^\gamma\ )$$

$$\vdots$$

$$\mathbf{c}_t^*(\theta^t) = \sigma(h^t, \theta^{t-1}, \theta_t)$$

$$\vdots$$

As with full information, a quota system is feasible only if it can be implemented by, in this case, a Perfect Bayesian equilibrium strategy pair (σ, μ). The pair (μ, σ) and a belief system $F_i^*(\theta_{it}|\ h^t, \tilde{\theta}^t),\ i = 1, \ldots n$ 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 (given in (8)) given usage history h^t, disclosure history $\tilde{\theta}^t$, i's current type θ_{it}, 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 $\tilde{\theta}^{t-1}$ and θ_{it}, and given the strategies of others; and (iii) beliefs F^* satisfy Bayes' Rule wherever possible.

By definition, if (σ, μ) is a truth-telling PBE that implements $\mathbf{c}^*$, then $V_i(h^t, \theta^t, \sigma, \mu|\ \theta_{it}) = U_i(\omega_t, c^*(\theta^t), \theta_{it})$ holds along the equilibrium path.

Note that at the disclosure stage, countries contemplate disclosure deviations from the from the prescribed plan, taking account of the fact that they have the freedom to deviate 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, a disclosure deviation, or a joint deviation in both use and disclosure — occurred.

The IA must recommend a quota system $\mathbf{c}^*$ and a PBE (σ, μ) that solves

$$\max_{\mathbf{c}^*} \sum_i \int_{\theta_0} U_i(\omega_0, c^*(\theta_0), \theta_{i0}) dF(\theta_0) \quad \text{subject to } \mathbf{c}^* \text{ implemented by the PBE } (\sigma, \mu). \tag{9}$$

Any quota that solves (9) is 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 $\mathbf{c}^*$ is fully compressed if for every country i, the quota c_{it}^* recommended to i at date t does not vary with the realized history of shocks θ^t. In other words, the quota is completely insensitive to the countries' realized preferences/production intensities for carbon.

To show this, we break down the problem in (9) into two parts, much as we did in the full information model. In Step 1, we solve a "relaxed planner's problem" that assumes, among other things, that carbon consumption can be imposed by the IA. Hence, only incentive constraints on information provision are imposed.

We characterize the solution to Step 1 in closed form, showing that it is, in fact, a fully compressed quota. Step 2 shows that the fully compressed solution can be implemented by a PBE in which the punishment following any deviation is itself compressed. In such a case, information need never be disclosed, and there is no learning about others' types on the equilibrium path or off.

4.3 Perfect Persistence

This Section takes up the special case where the shocks are perfectly persistent. The shock realized at $t = 0$ is realized once and for all. Formally, $\theta_0 = \theta_t$, for all t. This is a reasonable approximation when technological and environmental changes moves slowly relative to the frequency at which resource decisions are made. We later return to the general model.

Step 1 is to examine the solution to a "relaxed problem" as we did for full information. In the relaxed problem, the PBE constraint is replaced by a weaker constraint,

$$\int_{\theta_{-i}} U_i(\omega_0, \mathbf{c}^\star(\theta), \theta_i) dF_{-i}(\theta_{-i}) \;\geq\; \int_{\theta_{-i}} U_i(\omega_0, \mathbf{c}^\star(\tilde{\theta}_i, \theta_{-i}), \theta_i) dF_{-i}(\theta_{-i}) \quad \forall\, \tilde{\theta}_i \;\; \forall\, \forall\, i \tag{10}$$

With perfect persistence, there is only a single disclosure stage at the beginning of $t = 0$. Inequality (10) requires truth-telling at this stage. By requiring only (10), we ignore for now the compliance incentives. We will show that the optimal quota under the relaxed problem can be implemented by a full-blown truth-telling PBE.

Lemma 3 (Quota Compression) *With private shocks, the solution to the relaxed problem is a quota system $\mathbf{c}^\star$ characterized by a stationary usage profile $\mathbf{e}^\star$ that is independent of realized types. In particular, for each country i,*

$$e_i^\star \;=\; \frac{(1-\delta\gamma)\left[\int_{\underline{\theta}_i}^{\overline{\theta}} \theta_i dF_i(\theta_i)\right]}{n} \tag{11}$$

The proof is in the Appendix. Observe that the compressed rates $\mathbf{e}^*$ described by (11) obviously yield a compressed quota system. In this case, the corresponding quota system is given by

$$c_{it}^{*} \;=\; \frac{\int_{\underline{\theta}_i}^{\overline{\theta}} \theta_i dF_i(\theta_i)(1-\delta\gamma)}{n}\, \omega_0^{\gamma^t} \left(1 - \frac{\sum_i \int_{\underline{\theta}_i}^{\overline{\theta}} \theta_i dF_i(\theta_i)(1-\delta\gamma)}{n}\right)^{\gamma(1-\gamma^t)/(1-\gamma)} \tag{12}$$

Stark consequences. To take a particularly simple case, suppose that $F_i = F_j = F$, i.e., all countries draw their shocks from the same distribution. Then the optimal quota assigns them *identical* resource levels, despite the fact that the international agency can condition its recommendation on the information disclosed by each of the countries. If for instance, F is uniform on $[\underline{\theta}, \overline{\theta}]$, then the optimal quota is fully compressed and given explicitly by $e_i^{\star} = (1-\delta\gamma)(\overline{\theta}+\underline{\theta})/2n$. Notice, moreover, that the compressed quota is optimal even if there is very little noise in the distributions.

With private shocks, countries with realized usage values above the mean will generally extract less than under the full information optimum. Those below the mean will extract more. Generally, the informational rents accorded to low types gives them considerable "bargaining power." Relative to the full information optimum, high types subsidize low types.

To put this in concrete terms, fast-developing countries that end up 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).

Comparison with auctions. Given such stark consequences for any prospective climate agreement, the basic logic warrants some explanation (the formal proof is in the Appendix). A comparison with auctions is instructive. Given the structure of payoffs, it turns out that a country's payoff under any quota system is of the form $R_i(\theta_i) - \theta_i Q_i(\theta_i)$. Readers will immediately recognize the similarity to a firm in an oligopolistic industry, or to a seller in a procurement auction. The seller's cost parameter is θ_i, and with truthful disclosure it receives expected compensation $R_i(\theta_i)$ for a market share of $Q_i(\theta_i)$ units. Using standard monotonicity and envelope arguments, the expected payoff to a firm/country in an incentive compatible quota system (one satisfying (10)) can be expressed as

$$R_i(\overline{\theta}_i) - \overline{\theta}\; Q_i(\overline{\theta}_i) + \int_{\underline{\theta}_i}^{\overline{\theta}_i} \frac{F_i(\tilde{\theta}_i)}{f_i(\tilde{\theta}_i)} Q_i(\tilde{\theta}_i) f_i(\tilde{\theta}_i) d\tilde{\theta}_i$$

where Q_i must be weakly decreasing in the type θ_i.

In other words, in an incentive compatible quota system, a countries expected payoff is its payoff as the highest resource type $\overline{\theta}_i$, plus the usual "information rent" that depends only on the market share term Q_i. Critically, in most auction environments $Q_i(\theta_i)$ is either bounded or there are trade offs in the production quota between countries or between sellers and buyers. 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 in expectation or by a reduction in the likelihood of gains from trade with a buyer. 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.[15]

One exception to this auction logic is when the hazard rate F_i/f_i is strictly increasing. In that case, the gains from efficient sorting are 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 among the firms is fully compressed (which they refer to as "rigid") when hazard rates are increasing.[16] Roughly, a rigid/compressed production quota can then be shown to stochastically dominate any strictly decreasing one when weighted by an increasing hazard rate.[17]

In the present model, hazard rates need not be monotone. The difference between these results and our own concerns the description of Q_i. Using carbon as the leading example, $Q_i(\theta_i)$ in our model is *i's expected present value of stored carbon relative to its own resource consumption.* Formally, it is given by

$$Q_i(\theta_i) \ = \ \int_{\theta_{-i}} \sum_{t=0}^{\infty} \delta^t \log\left(\frac{1-\mathcal{E}_t(\theta)}{e_{it}(\theta)}\right) dF_{-i}(\theta_{-i})$$

The key point is that stored carbon is a public good. Hence, there are no allocative trade offs. All countries' "market shares" for carbon storage can be increased at once, and so the optimal mechanism should set all country's storage at its highest value $Q_i(\underline{\theta}_i)$ regardless of the realized type. For this reason, the compression occurs in the optimal quota even if the hazard rate is not monotone.

The Role of Commitment. Somewhat more opaque is the role of commitment. On the one hand, one might imagine the International Agency as a planner that literally solicits information on types. Since the quota is fully compressed, each country has no incentive to lie about its type. That being the case, however, a time consistent planner would use the information to implement the full information optimal quota after date $t = 1$. In turn, this destroys the initial incentive for truthful disclosure.[18] On the other hand, if there exists a PBE

[15]A recent paper by Pavan, Segal, and Toikka (2012) provides a comprehensive characterization of optimal mechanisms in Markov models of private information. Participants receive persistent, private shocks each period. Compliance with the planner's chosen mechanism is assumed, however, the participant's disclosure of their private information each period must be consistent with a PBE. Their characterization of the general monotonicity and envelope conditions required for optimal mechanisms is widely applicable across a large variety of dynamic models. We refer to the reader to their paper for further references.

[16]or if the maximum possible compensation from monopoly pricing is large enough.

[17]Their collusive mechanism allocates shares of production to firms who service an inelastic demand. The production shares must, on average, add to one.

[18]This point was originally made by Roberts (1984), in a cogent paper on optimal taxation in a dynamic "Mirleesian" economy.

consisting of punishment continuations that are *also* compressed, then truthful disclosure is not required to implement the optimal quota. If this is the case, then commitment is not necessary. This is precisely what we show below.

Implementing the relaxed planner's solution To implement the quota $\mathbf{e}^\star$ the IA can adapt the sanctions constructed in the full information case. That logic (outlined in Lemma 1) requires credible punishment at each counterfactual stage depends on even harsher punishment if the countries fail to carry out the sanction. This means, in turn, increasingly severe resource depletion must be threatened, itself made credible only by the threat of still further depletion later on, and so forth. The successive threats are only reached, of course, by further deviations at each counterfactual stage. The incentive constraints are therefore recursively defined.[19]

The difficulty with that construction is that a country may be able to manipulate the quota by first misreporting in the disclosure stage and then altering its resource utilization in the consumption stage. 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 varies across type, regardless of whether the prescription is on-path or off. Consequently, even though the quota $\mathbf{c}^\star$ itself is compressed, it may be difficult to implement in Perfect Bayesian equilibria if the sanctions are not. The result below by-passes this problem by showing that the optimal quota can be implemented by compressed sanctions.

Lemma 4 *With private shocks, the quota system $\mathbf{c}^\star$ that solves the relaxed planner's problem can be implemented by a Perfect Public Bayesian equilibrium (μ, σ) in which, for all countries i, and all histories h^t, i's usage strategy $\sigma_i(h^t, \tilde{\theta},\ \theta_i)$ does not vary in either the reported profile $\tilde{\theta}$ or its type θ_i.*

The proof, once again, is in the Appendix. Since the PBE does not, by definition, vary in the countries' private information, the punishments structure at each counterfactual layer τ is compressed, just as the optimal quota is.

Putting the two results, Lemma 3 and Lemma 4, together yields the following result.

Proposition 1 *With private, persistent shocks, the optimal quota $\mathbf{c}^\star$ is full compressed and, more specifically:*

(1) $\mathbf{c}^\star$ is characterized by fully compressed, stationary rates

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

[19] These are given by (19) in the proof of Lemma 1. See the Appendix for details.

for each country i, and

(2) $\mathbf{c}^\star$ *is implementable by a Perfect Public Bayesian equilibria* (μ, σ) *in which each country's prescribed usage rate after any history is stationary and compressed.*

4.4 Imperfectly Persistent Private Shocks

We now return to the general model with imperfectly persistent shocks. To start, consider Step 1 of our two-step solution method: the "relaxed" planner's problem:

$$\max_{\mathbf{c}^*} \sum_i \int_{\theta_0} U_i(\omega_0, c^*(\theta_0), \theta_{i0}) dF(\theta_0) \quad \text{subject to}$$

$$\int_{\theta_{-it}} U_i(\omega_t, c^*(\theta^t), \theta_{it}) dF_{-i}(\theta_{-it}|\theta_{t-1}) \geq \int_{\theta_{-it}} U_i(\omega_t, c^*(\theta^t \backslash \tilde{\theta}_{it}), \theta_{it}) dF_{-i}(\theta_{-it}|\theta_{t-1}) \;\; \forall\, \theta^t \;\forall\, \tilde{\theta}_{it} \;\forall\, i \;\forall\, t \tag{13}$$

This planner's problem is "super-relaxed" in the sense that it ignores consumption decisions and it ignores multi-period deviations in disclosure. Only incentive constraints for one-period deviations are considered. We proceed to show that the solution to this "super-relaxed" problem is a fully compressed quota, in which case the quota will also be immune from multi period deviations as well.

Lemma 5 (Quota compression with imperfectly persistent shocks). *With imperfectly persistent private shocks, the solution to the relaxed problem is a quota system* $\mathbf{c}^*$ *characterized by a non-stationary sequence of usage rates* $\{\mathbf{e}_t^*\}$ *such that each* $\mathbf{e}_t^*$ *is fully compressed (i.e., independent of realized types). In particular, for each country i, and each date t,*

$$e_{it}^* \;=\; \frac{(1-\delta\gamma)\int_{\underline{\theta}}^{\overline{\theta}} \hat{\theta}_{it} dF^t(\hat{\theta}_{it})}{n} \tag{14}$$

Notice that this quota is fully compressed but is not necessarily stationary insofar as it may vary over time due to drift in the distribution F^t. If, for instance, costs of climate change are expected to increase over time, then expected resource consumption types should drift downward (θ_i stochastically decreases over time). This means that the quota assigned to each country decreases which is what one should expect. The quota does not, however, vary with actual type realizations. Clearly this solution coincides with the perfect persistence model when types are perfectly correlated across time. In the case where the initial prior across all countries is the same, *the optimal quota yields no cross sectional variation, but does allow for time variation due to distributional drift.*

The basic logic of Lemma 5 is similar to the perfect persistence case. Once again, the payoff can be shown to be decomposed. As before, there is a “compensation” term and a “market share” term. There is also a third term which does not appear in the full persistence case. The third term comes from the fact that now, a country’s current disclosure decision can alter the beliefs of other countries about the disclosing country’s future type sequence. In turn, this can effect the original country’s quota. This term disappears whenever if the quota is compressed. Hence, we conjecture a solution that is fully compressed, then verify that it in fact solves a relaxed planner’s problem.

The issue of commitment. As with the perfect persistence case, commitment by the IA is critical. With full disclosure, the IA can use the conditional information at $t-1$ to make inferences on types at t. Clearly it would choose to use this information at each t. Since the sequential rationality of the quota is an additional constraint, the original (compressed) solution is no longer feasible, and so the countries are collectively worse off without the commitment.

Implementing the relaxed planner’s solution. The final step is to show PBE implementation using compressed punishments. This last step can be adapted from the perfect persistence case.

Lemma 6 *With private, imperfectly persistent shocks, the quota system* $\mathbf{c}^*$ *can be implemented by a Perfect Public Bayesian equilibrium* (μ, σ) *in which, for all countries* i*, and histories* h^t*, i's usage strategy* $\sigma_i(h^t, \tilde{\theta}^t,\ \theta_{it})$ *does not vary in the disclosure history* $\tilde{\theta}^t$ *or its type* θ_{it}*.*

Since the PBE does not, by definition, vary in the countries’ private information, the punishments structure at each counterfactual layer τ is compressed, just as the optimal quota is. The proof largely mimics the steps of the proof of Lemma 4, and so we omit the details.

The last two results imply our most general result, stated below.

Proposition 2 *With private, imperfectly persistent shocks, the IA’s optimal quota is given by a fully compressed quota* $\mathbf{c}^*$ *such that*

(1) $\mathbf{c}^*$ *is characterized by full compressed (non-stationary) rates*

$$e^*_{it} \;=\; \frac{(1-\delta\gamma)\int_{\underline{\theta}}^{\overline{\theta}} \hat{\theta}_{it} dF^t(\hat{\theta}_{it})}{n}$$

for each country i*, and*

(2) $\mathbf{c}^*$ *is implementable by a Perfect Public Bayesian equilibria* (μ, σ) *in which each country's prescribed usage rate after any history is compressed.*

5 Conclusion

This paper studies dynamic mechanisms for global commons with environmental externalities. Using carbon consumption is the leading example, we generalize the dynamic "fish war" game to allow for countries to have heterogeneous but direct benefits from resource conservation. We take up the case where countries incur serially correlated payoff/technology shocks. These shocks alter the way that countries evaluate the relative benefits and costs of carbon consumption over time.

In this context, an optimal quota system is an international climate agreement that assigns carbon 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 shows that the optimal quota system is fully compressed. The result is stark, as it suggests that the quota can be tailored only 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.

The quota compression requirement stands in firm contrast to the numerous international proposals (see Bodansky (2004) for a survey) advocating maximal flexibility in making adjustments particular characteristics of each country. Article 4 in the fundamental UN Framework Convention on Climate Change is fairly explicit about this. It makes repeated reference to 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 be reconciled to these approaches to the extent that some information about the shocks is publicly available. This would be true if, for instance, the data is gathering process occurs above and beyond the reach of sovereign filters. This may be reasonably presumed of the climate science itself. It is less reasonable, perhaps, when it comes to country-specific estimations of economic costs benefits since these rely to a large extent on each country's yearly disclosure of its national income accounts.

6 Appendix: Proofs of the Results

Proof of Lemma 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.

Before proceeding, observe that each country's payoff after history h^t under usage profile σ may be expressed as a recursive payoff

$$V_i(h^t, \sigma, \theta) \ = \ \theta_i \log \sigma_i(h^t,\ \theta) + (1-\theta_i) \log \ (\omega_t - \sum_{j=1}^{n} \sigma_j(h^t,\ \theta)) \ + \ \delta \ V_i(h^{t+1}, \sigma, \theta) \tag{15}$$

where decisions in period t determine the history h^{t+1} entering $t+1$.[20] If a usage profile σ implements a quota system $\mathbf{c}^*(\theta)$, then it follows from (2) and (15) that $V_i(h^0, \sigma, \theta) = U(\omega_0, \mathbf{c}^*(\theta), \theta_i)$.

Iterating on this payoff starting from ω_0, for any dynamic path $\mathbf{e}$ of usage rates, the long run payoff to a country i can be expressed as

$$\begin{aligned} & \frac{\omega_0}{1-\delta\gamma} \ + \ \sum_{t=0}^{\infty} \delta^t \left[\left(\frac{1}{1-\delta\gamma} - \theta_i \right) \log(1-\mathcal{E}_t) \ + \ \theta_i \log e_{it} \right] \\ \equiv \ & \frac{\omega_0}{1-\delta\gamma} \ + \ \sum_{t=0}^{\infty} \delta^t \ u_{it} \end{aligned} \tag{16}$$

Here, u_{it} captures the long run effect on payoffs of the extraction profile $\mathbf{e}_t$ chosen at date t. We refer to u as the *flow payoff* even though u_{it} includes future as well as present effects of the current profile $\mathbf{e}_t$. The critical feature used in the proof is the fact that each flow payoff is unbounded below. In the rest of the proof we make use of this notation and, moreover, drop the first term $\frac{\omega_0}{1-\delta\gamma}$ 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

$$V_i^*(\theta) \equiv \frac{\omega_0}{1-\delta\gamma} \ + \ \frac{1}{1-\delta} \left[\left(\frac{1}{1-\delta\gamma} - \theta_i \right) \log(1-\mathcal{E}^*(\theta)) \ + \ \theta_i \log e_i^*(\theta) \right] \tag{17}$$

to each country i.

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

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

Next, for $\tau \geq 1$, let $\mathbf{e}^{\tau-1}(\theta)$ be a stationary profile of usage rates and $V_i^{\tau-1}(\theta)$ the associated long run payoff for each country. We define the stationary profile $\mathbf{e}^{\tau}(\theta)$ and associated payoff $V_i^{\tau}(\theta)$ for each country as follows.

For each τ, all countries choose $\mathbf{e}^{\tau}(\theta)$ (to be defined shortly) for that period. This yields each country a flow payoff of $u_i^{\tau}(\theta)$. After one period, the IA carries out a randomization in which the system remains in the τ state with probability ρ. With probability $1-\rho$ each period the countries transition to the $\tau-1$ state which yields payoffs $V_i^{\tau-1}(\theta)$. The payoff $V_i^{\tau}(\theta)$ is therefore defined by

$$V_i^{\tau}(\theta) = u_i^{\tau}(\theta) + \delta[(1-\rho)V_i^{\tau-1}(\theta) + \rho V_i^{\tau}(\theta)]$$

Using the definition in (16) of an arbitrary flow payoff, the payoff in state τ is given by

$$V_i^{\tau}(\theta) = \frac{1}{1-\delta\rho}\left[\left(\frac{1}{1-\delta\gamma} - \theta_i\right)\log(1-\mathcal{E}^{\tau}(\theta)) \ + \ \theta_i \log e_i^{\tau}(\theta)\right] + \frac{\delta(1-\rho)}{1-\delta\rho}V_i^{\tau-1}(\theta) \tag{18}$$

To complete the recursive definition, we need to define $\mathbf{e}^{\tau}(\theta)$. This will be constructed to satisfy the incentive constraint in stage $\tau-1$. Specifically, if it turns out that a country deviates in state $\tau-1$ then the countries transition to the state τ in which usage rates are given by $\mathbf{e}^{\tau}(\theta)$ in the next period. Note that if there are no deviations, each state $\tau-1$ transitions to a lower state $\tau-2$ with probability ρ each period until finally play returns to the optimal quota $\mathbf{e}^{0}(\theta) \equiv \mathbf{e}^{*}(\theta)$. Consequently, for each $\tau \geq 1$, $\mathbf{e}^{\tau}(\theta)$ is defined to satisfy:

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

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

Using the definition of V_i^{τ} in (18), one can show that the incentive constraint holds if for all i,

$$\begin{aligned} V_i^{\tau-1}(\theta) \ \geq \ & \frac{1-\delta\rho}{1-\delta\rho-\delta^2(1-\rho)} \ \max_{e_i}\left[\left(\frac{1}{1-\delta\gamma} - \theta_i\right)\log(1-\mathcal{E}_{-i}^{\tau-1}(\theta)+e_i) \ + \ \theta_i \log e_i\right] \\ & + \ \frac{\delta^2}{1-\delta\rho-\delta^2(1-\rho)}\left[\left(\frac{1}{1-\delta\gamma} - \theta_i\right)\log(1-\mathcal{E}^{\tau}(\theta)) \ + \ \theta_i \log e_i^{\tau}(\theta)\right] \end{aligned} \tag{19}$$

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 $\{\mathbf{e}^{\tau}(\theta)\}_{\tau=0}^{\infty}$ is recursively constructed so that for each $\mathbf{e}^{\tau-1}(\theta)$, $\mathbf{e}^{\tau}(\theta)$ is chosen to satisfy these incentive constraints. Now let $h^t(e^{\tau}(\theta))$ denote the history at date t such that the last known deviation was in the profile $e^{\tau}(\theta)$. Then construct σ such that $\sigma_i(h^t(e^\tau(\theta)),\theta) = \omega_t e_i^{\tau+1}(\theta),\ t \geq 1$ and $\sigma_i(h^0,\theta) = \omega_0 e_i^*(\theta) \equiv c_i^*(\theta)$. By construction, the profile is subgame perfect, and so it implements $\mathbf{c}^*(\theta)$. ∎

Proof of Lemma 2. Observe that both the optimal rates $\mathbf{e}^*(\theta)$ and MPE rates $\mathbf{e}^\circ(\theta)$ are stationary. Recall that the long run payoffs for each, abbreviated here by $V^*(\omega_t,\theta)$ and $V^\circ(\omega_t,\theta)$, resp., take the form given by equation (17) (in the Proof of Lemma 1) where ω_t replaces ω_0 in the expression. In order for the optimal quota system to dominate the MPE fish war for country i, we require

$$
\begin{aligned}
& V_i^*(\omega_t,\theta) \;-\; V_i^\circ(\omega_t,\theta) \\
= \;& \left(\frac{1}{1-\delta\gamma} - \theta_i\right) \log(\frac{1-\mathcal{E}^*(\theta)}{1-\mathcal{E}^\circ(\theta)}) \;+\; \theta_i \log(\frac{e_i^*}{e_i^\circ}) \\
\geq \;& 0
\end{aligned}
$$

If θ_i is close enough to one, the inequality fails for $\delta\gamma$ close to zero However, it may also fail even in the limiting case where $\delta\gamma$ is close to one. Consider the limit

$$\lim_{\delta\gamma\to 1}\left[\left(\frac{1}{1-\delta\gamma} - \theta_i\right) \log(\frac{1-\mathcal{E}^*(\theta)}{1-\mathcal{E}^\circ(\theta)}) \;+\; \theta_i \log(\frac{e_i^*(\theta)}{e_i^\circ(\theta)})\right]$$

One can rewrite this limit as

$$\lim_{\delta\gamma\to 1}\left[\log\left(\frac{\frac{1}{1-\delta\gamma} - \mathcal{E}^\circ(\theta) + (\mathcal{E}^\circ(\theta) - \mathcal{E}^*(\theta))}{\frac{1}{1-\delta\gamma} - \mathcal{E}^\circ(\theta)}\right)^{\frac{1}{1-\delta\gamma}} \;+\; \log\left((\frac{e_i^*(\theta)}{e_i^\circ(\theta)})(\frac{1-\mathcal{E}^\circ(\theta)}{1-\mathcal{E}^*(\theta)})\right)^{\theta_i}\right]$$

Now recall that

$$e_i^\circ(\theta) = \frac{\theta_i(1-\delta\gamma)}{1+\Theta_{-i}(1-\delta\gamma)} \quad > \quad e_i^*(\theta) = \frac{\theta_i(1-\delta\gamma)}{n} \quad \forall\ i\ \ \forall\ \theta.$$

which follows from the fact that $\Theta_{-i}(1-\delta\gamma) < n-1$. In other words, The MPE extraction is always higher than that of the social optimum, regardless of the type profile of the countries. This is the standard free rider problem for the heterogeneous country case. This means, in turn, that $\mathcal{E}^\circ(\theta) - \mathcal{E}^*(\theta) > 0$, i..e, the aggregate rate is lower in the social optimum. Consequently, the first term in the bracket converges to $\frac{n-1}{n}\Theta$ as $\delta\gamma \to 1$.[21] The second term converges to $\log n^{-\theta_i}$. Hence, the inequality holds in the limit iff

$$\frac{\Theta}{n} \;\geq\; \theta_i \frac{\log n}{n-1} \tag{20}$$

[21]To see this, observe that $\left(\frac{\frac{1}{1-\delta\gamma} - \mathcal{E}^\circ(\theta) + (\mathcal{E}^\circ(\theta) - \mathcal{E}^*(\theta))}{\frac{1}{1-\delta\gamma} - \mathcal{E}^\circ(\theta)}\right)^{\frac{1}{1-\delta\gamma}} \to \exp(\mathcal{E}^\circ(\theta) - \mathcal{E}^*(\theta)) = \exp(\frac{n-1}{n}\Theta)$ as $\delta\gamma \to 1$

One can check that this always holds for the homogeneous case where $\theta_i = \theta_j$ for any pair of countries i and j. However, the constraint *fails* if a country's usage type is sufficiently large, relative to the average type $\frac{\Theta}{n}$.

This means that the fish war may not always be useful as a sanctioning threat in the optimal quota. But we can find conditions for which this is true for any candidate quota, even when the countries are exceedingly patient. This is now be shown.

Consider any quota system $\mathbf{c}^*(\theta)$. To show that $\mathbf{c}^*(\theta)$ can be implemented by an equilibrium profile σ^*, it suffices (via the one-shot deviation principle) to show on any on-path state ω_t^*,

$$V_i^*(\omega_t^*, \theta) \;\geq\; \max_{c_{it}} \;\left\{ \theta_i \log c_{it} + (1-\theta_i)\log(\omega_t - c_{it} - \sum_{j\neq i}^{n} c_j^*(\theta)) + \delta V_i^\circ(\omega_{t+1}, \theta) \right\}$$

where $\omega_{t+1} = (\omega_t - c_{it} - \sum_{j\neq i}^n c_j^*(\theta))^\gamma$. Expressing the inequality in terms of rates $\mathbf{e}^*(\theta)$ rather than consumption $\mathbf{c}^*(\theta)$, we rewrite the inequality above as

$$\max_{e_{it}} \;\left\{ \theta_i(\log e_{it} - \log e_i^*(\theta)) + (1-\theta_i)(\log(1 - e_{it} - \sum_{j\neq i}^{n} e_{jt}^*(\theta)) - \log(1 - \mathcal{E}^*(\theta)) \right\}$$

$$\leq\; \delta\left(\; V_i^*(\omega_{t+1}^*, \theta) - V_i^\circ(\omega_{t+1}, \theta)\;\right)$$

Using the form of payoff in (17), the incentive constraint can be expressed as

$$\max_{e_{it}} \;\left\{ \theta_i \log\frac{e_{it}}{e_i^*(\theta)} + (1-\theta_i)\log(\frac{1 - e_{it} - \sum_{j\neq i}^{n} e_{jt}^*(\theta)}{1-\mathcal{E}^*(\theta)}) + \frac{\delta\gamma}{1-\delta\gamma}\log(\frac{1 - e_{it} - \sum_{j\neq i}^{n} e_{jt}^*(\theta)}{1-\mathcal{E}^*(\theta)}) \right\}$$
$$\leq\; \frac{\delta}{1-\delta}\left[\left(\frac{1}{1-\delta\gamma} - \theta_i\right)\log(\frac{1-\mathcal{E}^*(\theta)}{1-\mathcal{E}^\circ(\theta)}) \;+\; \theta_i\log(\frac{e_i^*(\theta)}{e_i^\circ(\theta)})\right] \tag{21}$$

When $\mathbf{e}^*(\theta)$ is the unconstrained optimal quota then this incentive constraint will will not necessarily hold if δ is small. If δ is large, however, then one can verify that it does hold. To see this, notice, first that the left -hand side is bounded above for all $\delta\gamma$ (since in any best response e_{it} implies that $\log(\frac{1-e_{it}-\sum_{j\neq i}^n e_{jt}^*(\theta)}{1-\mathcal{E}^*(\theta)}) < 0$). Hence, the incentive constraint can be shown to hold for δ close to one. if term inside the square brackets $[\cdot]$ is positive. But that has already been shown in the limiting case where $\delta\gamma$ close to one whenever $\frac{\Theta}{n} \geq \theta_i\frac{\log n}{n-1}$ which is the inequality given in (20).

Evaluating (20) In the worst case, a country has the highest extraction type $\theta_i = \overline{\theta}$ and faces the rest of the world consisting of countries that are low extraction (" conservation-oriented") types $\theta_j = \underline{\theta}$ for $j \neq i$. That "worst case scenario" constitutes the inequality (6) in the Lemma and, if satisfied, ensures that the optimal quota system is sustainable by a threat of a fish war. ∎

Proof of Lemma 3. Let $\mathbf{c}^\star$ denote a candidate optimal quota system induced by the relaxed planner's problem, and let $\mathbf{e}$ denote the rates corresponding to $\mathbf{c}^*$. Notice that if (σ, μ) satisfies the truth-telling constraint (10), then any arbitrary report $\tilde{\theta}_j$ by country j should be believed to be j's true type with probability one by country i. Consequently, if profile $\tilde{\theta}$ is disclosed by all the countries, then a country of type θ_i has a long-run payoff under a quota (in rates) $\mathbf{e}$ given by,

$$\begin{aligned}
U_i(\omega_0, \mathbf{c}^*(\tilde{\theta}), \theta) &= \sum_{t=0}^{\infty} \delta^t \left[\theta_i \log \omega_t e_{it}(\tilde{\theta}) + (1-\theta_i) \log \omega_t (1 - \mathcal{E}_t(\tilde{\theta}))\right] \\
&= \sum_{t=0}^{\infty} \delta^t \left[\log \omega_t + \theta_i \log e_{it}(\tilde{\theta}) + (1-\theta_i) \log (1 - \mathcal{E}_t(\tilde{\theta}))\right] \\
&= \sum_{t=0}^{\infty} \delta^t \log \left(\omega_0^{\gamma^t} \prod_{j=0}^{t-1} (1 - \mathcal{E}_j(\tilde{\theta}))\right) \; + \; \theta_i \sum_{t=0}^{\infty} \delta^t \log e_{it}(\tilde{\theta}) \\
&\qquad + \; (1-\theta_i) \sum_{t=0}^{\infty} \delta^t \log (1 - \mathcal{E}_t(\tilde{\theta})) \\
&= \frac{\log \omega_0}{1-\delta\gamma} + \sum_{t=0}^{\infty} \delta^t \sum_{j=0}^{t} \gamma^{t-j} \log(1 - \mathcal{E}_j(\tilde{\theta})) \; - \; \theta_i \sum_{t=0}^{\infty} \delta^t \log \left(\frac{1 - \mathcal{E}_t(\tilde{\theta})}{e_{it}(\tilde{\theta})}\right) \\
&\equiv r_i(\tilde{\theta}) \; - \; \theta_i \; q_i(\tilde{\theta})
\end{aligned}$$

According to this definition, for any type profile θ,

$$r_i(\theta) \equiv \frac{\log \omega_0}{1-\delta\gamma} + \sum_{t=0}^{\infty} \delta^t \sum_{j=0}^{t} \gamma^{t-j} \log(1 - \mathcal{E}_j(\theta)) \;\; \textit{and} \tag{22}$$

$$q_i(\theta) \equiv \sum_{t=0}^{\infty} \delta^t \log \left(\frac{1 - \mathcal{E}_t(\theta)}{e_{it}(\theta)}\right). \tag{23}$$

The payoff has a simple interpretation. $r_i(\theta)$ is the long run value of resource conservation given type profile θ reported to the IA. This value can be viewed as benefit from avoiding the costs of usage, e.g., the costs of climate change. By contrast, $q_i(\theta)$ 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. The interim values of these are given by

$$R_i(\theta_i) \;\equiv\; \int_{\theta_{-i}} r_i(\theta)dF_{-i}(\theta_{-i}) \quad \textit{and} \quad Q_i(\theta_i) \;\equiv\; \int_{\theta_{-i}} q_i(\theta)dF_{-i}(\theta_{-i}) \tag{24}$$

Hence, by definition,

$$\int V_i(h^0, \theta, \sigma |\, \theta_i)dF_{-i}(\theta_{-i}) \;\equiv\; R_i(\theta_i) \;-\; \theta_i\; Q_i(\theta_i) \tag{25}$$

The truth-telling constraint (10) may therefore be expressed as

$$R_i(\theta_i) \;-\; \theta_i\; Q_i(\theta_i) \;\geq\; R_i(\tilde{\theta}_i) \;-\; \theta_i\; Q_i(\tilde{\theta}_i) \quad \forall\, \theta_i \;\; \forall\, \tilde{\theta}_i \tag{26}$$

Standard arguments show that any R_i and Q_i will satisfy (26) iff Q_i is weakly decreasing and

$$R_i(\theta_i) \;-\; \theta_i\; Q_i(\theta_i) \;=\; R_i(\overline{\theta}_i) - \overline{\theta}\; Q_i(\overline{\theta}_i) + \int_{\theta_i}^{\overline{\theta}_i} Q_i(\tilde{\theta}_i)d\tilde{\theta}_i \tag{27}$$

A simple integration by parts argument applied to (27) yields

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

Hence, we may now write the planner's relaxed problem as

$$\max_{R_i(\overline{\theta}_i), Q_i(\cdot)} \quad \sum_i \left[R_i(\overline{\theta}_i) - \overline{\theta}\; Q_i(\overline{\theta}_i) + \int_{\underline{\theta}_i}^{\overline{\theta}_i} \frac{F_i(\tilde{\theta}_i)}{f_i(\tilde{\theta}_i)} Q_i(\tilde{\theta}_i) f_i(\tilde{\theta}_i) d\tilde{\theta}_i \right] \tag{28}$$

subject to Q_i weakly decreasing and (23), (22), (24), and (27).

Now consider a putative solution $(R_i(\overline{\theta}_i), Q_i(\cdot))$ to (28). Because Q_i is weakly decreasing, the value of (28) is bounded above by

$$\sum_i \left[R_i(\overline{\theta}_i) - \overline{\theta}\; Q_i(\overline{\theta}_i) + Q_i(\underline{\theta}_i) \int_{\underline{\theta}_i}^{\overline{\theta}_i} F_i(\tilde{\theta}_i) d\tilde{\theta}_i \right] \tag{29}$$

Provided that the value of (29) is finite, this upper bound can be achieved exactly by setting Q_i and R_i so that $Q_i(\theta_i) = Q_i(\underline{\theta}_i)$ for all $\theta_i < \overline{\theta}_i$.[22] Since we do not violate the constraints (22), (23), (24), and (27), then any proposed solution $(R_i(\overline{\theta}_i), Q_i(\cdot))$ in which Q_i was *not* constant on $[\underline{\theta}_i, \overline{\theta}_i)$ could not, in fact, be a solution to (28).[23]

[22] Note that the incentive constraints allow for a potential discontinuity of Q_i at $\overline{\theta}_i$.

[23] A critical part of the argument is that Q_i can be increased for all individuals and for all types simultaneously. This is because an increase in Q_i is acheived by *decreasing* resource consumption for everyone at once.

Hence, we set $Q_i(\theta_i) = Q_i(\underline{\theta}_i)$ for all $\theta_i < \overline{\theta}_i$ and proceed to show that (29) is finite. First, we show that $e_i(\theta)$ must be constant in θ given $Q_i(\theta_i) = Q_i(\underline{\theta}_i)$. To see why, observe first that (23) and (24) are discounted sums of concave functions and so maximizing $\sum_i Q_i(\underline{\theta}_i)\int_{\underline{\theta}_i}^{\overline{\theta}_i} F_i(\tilde{\theta}_i)d\tilde{\theta}_i$ as required by (29) implies that resource usage is stationary. That is, $e_{it}(\theta) = e_i(\theta)$ for all i and t. As a consequence,

$$Q_i(\underline{\theta}_i) = \frac{1}{1-\delta}\int_{\theta_{-i}} \log(\frac{1-\mathcal{E}(\theta_{-i},\underline{\theta}_i)}{e_i(\theta_{-i},\underline{\theta}_i)})dF_{-i}$$

Second, observe that by Jensen's Inequality,

$$Q_i(\underline{\theta}_i) \leq \frac{1}{1-\delta}\log(\int_{\theta_{-i}} \frac{1-\mathcal{E}(\theta_{-i},\underline{\theta}_i)}{e_i(\theta_{-i},\underline{\theta}_i)}dF_{-i})$$

and this inequality is strict unless $\frac{1-\mathcal{E}(\theta_{-i},\underline{\theta}_i)}{e_i(\theta_{-i},\underline{\theta}_i)}$ is constant across all θ_{-i}. Since the putative planner's problem (29) requires that one maximize $Q_i(\underline{\theta}_i)$, it follows that $\frac{1-\mathcal{E}(\theta_{-i},\underline{\theta}_i)}{e_i(\theta_{-i},\underline{\theta}_i)}$ must indeed be constant across θ_{-i}. Set $k_i = \frac{1-\mathcal{E}(\theta_{-i},\underline{\theta}_i)}{e_i(\theta_{-i},\underline{\theta}_i)}$. We argue that this, in turn, implies that $e_j(\theta)$ is constant in θ for all j. Summing over all $e_i(\theta)$ as $i = 1,\ldots,n$, one obtains $\mathcal{E}(\theta) = \mathcal{E}^\star = \frac{\sum_i \frac{1}{2k_i}}{1+\sum_i \frac{1}{2k_i}}$ which clearly does not vary with θ. We then have $\frac{1-\mathcal{E}^\star}{e_i^\star(\theta)} = k_i$ and so $e_i(\theta)$ cannot vary with θ.

It therefore suffices to look for a triple: $R_i(\overline{\theta}_i)$, $Q_i(\overline{\theta}_i)$, and $Q_i(\underline{\theta}_i)$ that maximizes (29). To find the solution, first evaluate (27) at $\theta_i = \underline{\theta}_i$ and set $Q_i(\theta_i) = Q_i(\underline{\theta}_i)$. Equation (27) then reduces to:

$$R_i(\underline{\theta}_i) - Q_i(\underline{\theta}_i)\overline{\theta}_i \;=\; R_i(\overline{\theta}_i) - \overline{\theta}\;Q_i(\overline{\theta}_i)$$

Substituting this back into (29) yields

$$\sum_i \left[R_i(\underline{\theta}_i) \;-\; Q_i(\underline{\theta}_i)\left(\overline{\theta}_i - \int_{\underline{\theta}_i}^{\overline{\theta}_i} F_i(\tilde{\theta}_i)d\tilde{\theta}\right)\right] \tag{30}$$

Notice in (30) above that i's type is $\theta_i = \overline{\theta} - \int_{\underline{\theta}_i}^{\overline{\theta}} F(\tilde{\theta}_i)d\tilde{\theta}$ which, by a simple integration by parts is equal to $\int_{\underline{\theta}_i}^{\overline{\theta}} \tilde{\theta}_i dF_i(\tilde{\theta}_i)$, the average type for country i. We use this fact together with the fact that $e_i(\theta)$ does not vary with type, all to rewrite the objective in (28) as

$$\max_{r_i(\underline{\theta}),q_i(\underline{\theta})} \;\sum_i \left[r_i(\underline{\theta}) \;-\; q_i(\underline{\theta})\left(\int_{\underline{\theta}_i}^{\overline{\theta}} \tilde{\theta}_i dF_i(\tilde{\theta}_i)\right)\right] \tag{31}$$

subject to (22) and (23). Since the functions $r_i(\underline{\theta})$ and $q_i(\underline{\theta})$ are simply choice variables of the planner, the solution to (37) yields the full information optimum in which each individual's type is $\int \theta_i dF_i$. Looking back to last section, it easy to check that the solution is stationary and symmetric and given by

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

where $\mathbf{e}^\star$ expresses the stationary solution in this information constrained problem (without usage incentives). This concludes the proof of Lemma 3. ∎

Proof of Lemma 4 The proof largely mimics the steps of Lemma 1 which we do not repeated 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,

$$\begin{aligned} V_i^{\tau-1} \;\geq\; \max_{\theta_i} \Bigg\{ & \frac{1-\delta\rho}{1-\delta\rho-\delta^2(1-\rho)} \;\max_{e_i} \left[\left(\frac{1}{1-\delta\gamma} - \theta_i\right) \log(1-\mathcal{E}_{-i}^{\tau-1}+e_i) \;+\; \theta_i \log e_i \right] \\ + \;& \frac{\delta^2}{1-\delta\rho-\delta^2(1-\rho)} \left[\left(\frac{1}{1-\delta\gamma} - \theta_i\right) \log(1-\mathcal{E}^\tau) \;+\; \theta_i \log e_i^\tau \right] \Bigg\} \end{aligned} \tag{32}$$

Hence, the difference between (32) and the analogous constraint (19) 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 (32) 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 $\sigma_i(h^0, , \tilde{\theta}, \theta_i) = \omega_0 e_i^\star$, 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 $\mathbf{e}^*$. ∎

Proof of Lemma 5. For any quota $\mathbf{c}^*$, let $\mathbf{e}$ denote the corresponding profile of extraction rates, where $\mathbf{c}_t^*(\theta^t) = \omega_t^*(\theta^{t-1})\mathbf{e}_t(\theta^t)$.

Using the expression for payoffs in (7), a country i of type θ_{it} that reports $\tilde{\theta}_{it}$ has a payoff on path given by

$$\int_{\theta_{-it}} U_i(\omega_t, c^*(\theta^t \backslash \tilde{\theta}_{it}), \theta_{it}) dF_{-i}(\theta_{-it}|\theta_{t-1}) = \mathcal{R}_{it}(\tilde{\theta}_{it}|\theta^{t-1}) - \theta_{it}\mathcal{Q}_{it}(\tilde{\theta}_{it}|\theta^{t-1}) + \mathcal{P}_{it}(\tilde{\theta}_{it}, \theta_{it}|\theta^{t-1}) \tag{33}$$

where

$$\mathcal{R}_{it}(\tilde{\theta}_{it}|\theta^{t-1}) \;\equiv\; \frac{\log \omega_t}{1-\delta\gamma} \;+\; \frac{1}{1-\delta\gamma}\int_{\theta^t} \log(1-\mathcal{E}_t(\theta^t \backslash \tilde{\theta}_{it})) dF_{-i}(\theta_{-it}|\theta_{t-1}) \tag{34}$$

$$\mathcal{Q}_{it}(\tilde{\theta}_{it}|\theta^{t-1}) \equiv \int_{\theta_{-it}} \log\left(\frac{1-\mathcal{E}_t(\theta^t \backslash \tilde{\theta}_{it})}{e_{it}(\theta^t \backslash \tilde{\theta}_{it})}\right) dF_{-i}(\theta_{-it}|\theta_{t-1}) \tag{35}$$

and

$$\begin{aligned}\mathcal{P}_{it}(\tilde{\theta}_{it},\theta_{it}|\theta^{t-1}) \equiv & \sum_{\tau=t+1}^{\infty} \delta^{\tau-t}\Bigg\{ \left[\sum_{j=1}^{\tau} \gamma^{\tau-j} \int_{\theta^j} \log(1-\mathcal{E}_j(\theta^j \backslash \tilde{\theta}_{it})) dF_{-i}^{j-t+1}(\theta_{-i}^j|\theta^{t-1}) dF_i^{j-t}(\theta_i^j|\theta_{it})\right] \\ & - \left[\int_{\theta^\tau} \theta_{i\tau} \log\left(\frac{1-\mathcal{E}_\tau(\theta^\tau \backslash \tilde{\theta}_{it})}{e_{i\tau}(\theta^\tau \backslash \tilde{\theta}_{it})}\right) dF_{-i}^{\tau-t+1}(\theta_{-i}^\tau|\theta_{t-1}) dF_i^{\tau-t}(\theta_i^\tau|\theta_{it})\right] \Bigg\}\end{aligned} \tag{36}$$

This payoff expression in (33) is calculated to consider the one-shot incentive to deviate given past and future disclosure is truthful.[24] We have as a relaxed program constraint,

$$\begin{aligned}& \mathcal{R}_{it}(\theta_{it}|\theta^{t-1}) - \theta_{it}\mathcal{Q}_{it}(\theta_{it}|\theta^{t-1}) + \mathcal{P}_{it}(\theta_{it},\theta_{it}|\theta^{t-1}) \\ \geq\ & \mathcal{R}_{it}(\tilde{\theta}_{it}|\theta_{t-1}) - \theta_{i0}\mathcal{Q}_{it}(\tilde{\theta}_{it}|\theta_{t-1}) + \mathcal{P}_{it}(\tilde{\theta}_{it},\theta_{it}|\theta^{t-1}) \qquad \forall\ \theta_{it}\ \forall\ \tilde{\theta}_{it}\ \forall\ \theta^{t-1}\ \forall\ i\ \forall\ t\end{aligned} \tag{37}$$

This constraint ignores consumption incentives, and focuses purely on one-shot disclosure incentives.

The relaxed problem can now be restated as:

$$\max_{\mathbf{c}^*} \sum_i \int_{\theta_0} U_i(\omega_0, c^*(\theta_0), \theta_{i0}) dF(\theta_0) \quad \text{subject to (37).} \tag{38}$$

The relaxed IC again implies a generalization of the monotonicity condition $\mathcal{Q}_{it}(\theta_{it}|\theta^{t-1}) - D_2\mathcal{P}_{it}(\theta_{it},\theta_{it}|\theta^{t-1}) \leq \mathcal{Q}_{it}(\tilde{\theta}_{it}|\theta^{t-1}) - D_2\mathcal{P}_{it}(\tilde{\theta}_{it},\theta_{it}|\theta^{t-1})$ whenever $\theta_{it} \geq \tilde{\theta}_{it}$, where $D_2\mathcal{P}_{it}$ denotes the partial derivative in its second argument.

In addition to monotonicity and a standard envelope condition holds, in which i's (ex ante)

[24] Note that in the case of perfect persistence studied in the previous section, $\mathcal{Q}_{it}(\tilde{\theta}_{it}|\theta^{t-1}) = Q_i(\tilde{\theta}_i)$.

long run expected payoff of the quota, beginning in date t is given by

$$\int_{\theta_{-it}} U_i(\omega_t, c^*(\theta^t), \theta_{it}) dF_{-i}(\theta_{-it}|\theta_{t-1}) \;=\; \mathcal{R}_{it}(\overline{\theta}_{it}|\theta^{t-1}) - \overline{\theta}_{it}\mathcal{Q}_{it}(\overline{\theta}_{it}|\theta^{t-1}) \;+\; \mathcal{P}_{it}(\overline{\theta}_{it}, \overline{\theta}_{it}|\theta^{t-1})$$
$$+ \int_{\theta_{it}}^{\overline{\theta}_{it}} \left(\mathcal{Q}_{it}(\hat{\theta}_i|\theta^{t-1}) - D_2\mathcal{P}_{it}(\hat{\theta}_i, \hat{\theta}_i|\theta^{t-1})\right) d\hat{\theta}_i \tag{39}$$

Using an integration by parts argument (as in the perfect persistence case),

$$\int_{\theta_t} U_i(\omega_t, c^*(\theta^t), \theta_t) dF^t(\theta^t) \;=\; \int_{\theta^{t-1}} \left[\mathcal{R}_{it}(\overline{\theta}_{it}|\theta^{t-1}) - \overline{\theta}_{it}\mathcal{Q}_{it}(\overline{\theta}_{it}|\theta^{t-1}) + \mathcal{P}_{it}(\overline{\theta}_{it}, \overline{\theta}_{it}|\theta^{t-1})\right] dF^{t-1}(\theta^{t-1})$$
$$+ \int_{\theta^{t-1}} \left[\int_{\underline{\theta}_{it}}^{\overline{\theta}_{it}} \left(\mathcal{Q}_{it}(\theta_{it}|\theta^{t-1}) - D_2\mathcal{P}_{it}(\theta_{it}, \theta_{it}|\theta^{t-1})\right) F_i(\theta_{it}|\theta_{t-1}) d\theta_{it}\right] dF^{t-1}(\theta^{t-1})$$
$$\equiv\; \mathcal{R}^*_{it}(\overline{\theta}_{it}) - \overline{\theta}_{it}\mathcal{Q}^*_{it}(\overline{\theta}_{it}) + \mathcal{P}^*_{it}(\overline{\theta}_{it}, \overline{\theta}_{it}) \;+\; \int_{\underline{\theta}_{it}}^{\overline{\theta}_{it}} \left(\mathcal{Q}^*_{it}(\theta_{it}) - D_2\mathcal{P}^*_{it}(\theta_{it}, \theta_{it})\right) F^t_i(\theta_{it}) d\theta_{it}$$

The planner's ex ante criterion from date t onward can now be stated as

$$\max \sum_i \left\{\mathcal{R}^*_{it}(\overline{\theta}_{it}) - \overline{\theta}_{it}\mathcal{Q}^*_{it}(\overline{\theta}_{it}) + \mathcal{P}^*_{it}(\overline{\theta}_{it}, \overline{\theta}_{it}) \;+\; \int_{\underline{\theta}_{it}}^{\overline{\theta}_{it}} \left(\mathcal{Q}^*_{it}(\theta_{it}) - D_2\mathcal{P}^*_{it}(\theta_{it}, \theta_{it})\right) F^t_i(\theta_{it}) d\theta_{it}\right\}$$

subject to the monotonicity condition on $Q^*_{it} - D_2 P^*_{it}$.

As in the perfect persistence case, the monotonicity condition implies that the planner's criterion from date t is maximized by setting $\mathcal{Q}^*_{it}(\theta_i) = \mathcal{Q}^*_{it}(\underline{\theta}_{it})$ and $D_2\mathcal{P}^*_{it}(\theta_{it}, \theta_{it}) = D_2\mathcal{P}^*_{it}(\underline{\theta}_{it}, \underline{\theta}_{it})$ for all $\theta_i < \overline{\theta}_i$ for all countries i. The planner's criterion becomes The planner's ex ante criterion from date t onward can now be stated as

$$\max \sum_i \left\{\mathcal{R}^*_{it}(\overline{\theta}_{it}) - \overline{\theta}_{it}\mathcal{Q}^*_{it}(\overline{\theta}_{it}) + \mathcal{P}^*_{it}(\overline{\theta}_{it}, \overline{\theta}_{it}) \;+\; \left(\mathcal{Q}^*_{it}(\underline{\theta}_{it}) - D_2\mathcal{P}^*_{it}(\underline{\theta}_{it}, \underline{\theta}_{it})\right) \int_{\underline{\theta}_{it}}^{\overline{\theta}_{it}} F^t_i(\theta_{it}) d\theta_{it}\right\} \tag{40}$$

We now argue that a solution to this problem is a quota at date t, $\mathbf{e}^*_t$ that is fully compressed. The argument largely repeats the steps of the perfect persistence case in Lemma 3. The argument is as follows. First, observe that by the definitions in (34), (35), and (36), the value of the date t quota $\mathbf{e}^*_t$ only appears in $\mathcal{R}^*_{it}$ and $\mathcal{Q}^*_{it}$. In particular, it follows from Jensen's inequality that

$$\mathcal{Q}^*_{it}(\underline{\theta}_{it}) \leq \log\left(\int_{\theta^{t-1}} \int_{\theta_{-it}} \frac{1 - \mathcal{E}_t(\theta^t \backslash \underline{\theta}_{it})}{e_{it}(\theta^t \backslash \underline{\theta}_{it})} dF_{-i}(\theta_{-it}|\theta_{t-1}) dF^{t-1}(\theta^{t-1})\right) \tag{41}$$

with strict inequality whenever $\frac{1-\mathcal{E}_t(\theta^t \backslash \underline{\theta}_{it})}{e_{it}(\theta^t \backslash \underline{\theta}_{it})}$ is constant across all θ^{t-1} and across all θ_{-i}. This implies that $\mathbf{e}_t^*$ is constant over these variables, in other words, $\mathbf{e}_t^*$ is fully compressed. Since the planner's problem in (40) is indeed maximized when (41) is an equality, the date t quota is compressed.

Now, by Blackwell's Principle, the solutions for $\mathbf{e}_\tau^*$ for $\tau > t$ that the planner chooses at date t should be consistent with those that we would choose at those future date date τ. This is verified by observing that when all future $\mathbf{e}_\tau^*$ chosen in those dates are compressed, it follows that $D_2\mathcal{P}_{it}^*(\underline{\theta}_{it}, \underline{\theta}_{it}) = 0$.[25] Hence, using this Envelope condition, the planner's Bellman's equation at date t only varies with the date-t quota $\mathbf{e}_t^*$ and not future quotas.

We can now use this fact and extreme compression to characterize the solution to the planner's problem. Using the same techniques as the perfect persistent case,[26] the planner's problem may be restated as: choose $\mathcal{R}_{it}^*(\underline{\theta}_{it})$, $\mathcal{Q}_{it}^*(\underline{\theta}_{it})$, and $\mathcal{P}_{it}^*(\underline{\theta}_{it}, \underline{\theta}_{it})$, to solve the planner's problem at t,

$$\max \sum_i \left\{ \mathcal{R}_{it}^*(\underline{\theta}_{it}) \ + \ \mathcal{P}_{it}^*(\underline{\theta}_{it}, \underline{\theta}_{it}) \ - \ \mathcal{Q}_{it}^*(\underline{\theta}_{it}) \int_{\theta_{it}} \theta_{it} dF_i^t(\theta_{it}) d\theta_{it} \right\} \tag{42}$$

Notice that this is the value of the full information planner's problem when each country's type is $\theta_{it} = \int_{\theta_{it}} \theta_{it} dF_i^t(\theta_{it}) d\theta_{it}$ in each period t. Given compression, the solution to this problem yields the same solution as the full information model with public shocks. One can verify that the solution is

$$e_{it}^* \ = \ \frac{(1-\delta\gamma)\int_{\underline{\theta}}^{\overline{\theta}} \hat{\theta}_{it} dF^t(\hat{\theta}_{it})}{n}$$

as required in the Lemma. ∎

References

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

[25] This follows from the fact that $\mathcal{P}_{it}^*$ is of the form $\int g(x)f(x|y)dx$, where $g(x)$ is a real valued function and $f(x|y)$ a regular conditional C^1 density with full support. Then $D_2\mathcal{P}_{it}^*(\underline{\theta}_{it}, \underline{\theta}_{it}) \equiv g(\underline{x}) \int \frac{\partial f(x|y)}{\partial y} dx = 0$.

[26] Briefly, evaluating i's conditional payoffs from date t onward when he is type $\theta_{it} = \underline{\theta}_i$ and combining it with the envelope- adjusted payoff gives,

$$\mathcal{R}_{it}^*(\underline{\theta}_{it}) - \underline{\theta}_{it}\mathcal{Q}_{it}^*(\underline{\theta}_{it}) + \mathcal{P}_{it}^*(\underline{\theta}_{it}, \underline{\theta}_{it}) = \mathcal{R}_{it}^*(\overline{\theta}_{it}) - \overline{\theta}_{it}\mathcal{Q}_{it}^*(\overline{\theta}_{it}) + \mathcal{P}_{it}^*(\overline{\theta}_{it}, \overline{\theta}_{it}) + \int_{\underline{\theta}_{it}}^{\overline{\theta}_{it}} \left(\mathcal{Q}_{it}(\theta_{it}|\theta^{t-1}) - D_2\mathcal{P}_{it}(\theta_{it}, \theta_{it}|\theta^{t-1}) \right) d\theta_{it}$$

Then using the monotonicity condition, we set $\mathcal{Q}_{it}^*$ and $\mathcal{P}_{it}^*$ at $\theta_{it} = \underline{\theta}_{it}$ in the integral, re-arrange terms, and substitute it into the planner's criterion (40). Finally, setting $D_2\mathcal{P}_{it}(\theta_{it}, \theta_{it}|\theta^{t-1}) = 0$ to account for compression at all future dates, we obtain the planner's criterion in (42).

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

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

[4] 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.

[5] Acemoglu, D. and J. Robinson (2008), "Persistence of Powers, Elites and Institutions," *American Economic Review*, 98: 267-293.

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

[7] 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.

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

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

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

[11] 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.

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

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

[14] Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski (2011), "Optimal Taxes on Fossil Fuel in General Equilibrium," NBER working paper.

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

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

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

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

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

[20] Pavan, A., I. Segal, and J. Toikka (2012), "Dynamic Mechanism Design," Stanford University mimeo.

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

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

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

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

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