

Dynamic Mechanism Design for a Global Commons

Rodrigo Harrison* and Roger Lagunoff†

September 15, 2013‡

ABSTRACT

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

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

Key Words and Phrases: Dynamic mechanism design, global commons, climate change, optimal quota, full compression, fish wars, Perfect Bayesian equilibria, international agency.

*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. +1-202-687-1510, lagunofr@georgetown.edu, www.georgetown.edu/lagunoff/lagunoff.htm.

‡We thank seminar participants at NYU, CETC Montreal, the NBER, USACH (Chile), and especially Sandeep Baliga for valuable comments and suggestions.

1 Introduction

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

Examples include deforestation, species extinction, and ozone depletion. A leading example is the use of carbon based resources, resulting in increased concentrations of green house gases (GHG). These gases, in turn, 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 that address 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.² 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.

Notwithstanding some earlier contributions by 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 may be that global commons environments often combine a common pool problem with a global environmental externality. This makes for a challenging design problem. The global scale of GHG emissions, for instance, requires that most countries be involved in the negotiations. Moreover, the accumulation of atmospheric CO₂ is an inherently dynamic process. 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

¹IPCC Fourth Assessment Report: Climate Change 2007.

²See Arava, et. al. (2010) for a summary of the mechanisms in place and their rationales.

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

catastrophic losses and discount rates. Aggregate estimates of costs mask significant differences in impacts across sectors, regions and populations and very likely underestimate damage costs because they cannot include many non-quantifiable impacts...” (IPCC 2007 Synthesis Report).

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

To address the issues in a tractable way, the present paper posits an infinite horizon model of global 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 because it allows the country to avoid the environmental costs of global resource consumption. In the model, the conservation benefits are assumed to be heterogeneous across countries and 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, demographic, and politico-economic influences. The shock process in our model captures a common feature in many commons problems: environmental costs are difficult to forecast and often vary widely across countries.⁴

The model generalizes the classic common pool resource model of Levhari and Mirman (1980).⁵ In their model, Levhari and Mirman study the strategic allocation of a depletable, open access resource such as fisheries 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 their model only for instrumental reasons: preserving the stock allows one to smooth consumption.

This paper modifies the common pool 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.

⁴Burke, et. al. (2011) find, for example, widely varying estimates of the effect of climate change on US agriculture when climate model uncertainty is taken into account. Desmet and Rossi-Hansberg (2013) quantify cross country variation in a calibrated model of spatial differences in welfare losses across countries due to global warming. These differences come primarily from geography but may be amplified by trade frictions, migration, and energy policy.

⁵See Long (2011) for a survey of the vast literature since the 1980 paper.

A *quota system* in this context 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 incentive-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 past disclosure and usage history. Implementability in PBE is a natural restriction as it requires the quota system to be *dynamically self-enforcing* in both compliance and in disclosure of information.

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.

We first analyze the case of full information as a benchmark. Under full information, each country’s payoff for consumption and conservation is known and there are no shocks. The first-best 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. We show that the first-best quota is implementable under full information. Significantly, the implementability does not depend on discounting. Because the effects of full depletion are catastrophic, cheating is deterred by graduated punishments that further deplete the stock each time a country violates its prescribed resource use along the punishment path.

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

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

To illustrate what full compression means, consider two countries that draw their types from the same distribution for carbon use, ex ante. Suppose the realized shocks at some point are such that one country ends up with high resource needs and/or low costs to climate change, while the other ends up with low resource needs and/or high environmental costs. Full compression then implies that the same quota is assigned to each.

For expositional reasons, we first establish full compression in a special case of the general model in which the shocks are perfectly persistent; each country privately

draws its realized payoff type once and for all at date $t = 0$. If countries are sufficiently patient, then the optimal quota is shown to be both fully compressed and time-stationary. We then extend the result to the case of imperfectly persistent shocks that follow a martingale. Once again, the quota exhibits full compression.

The compression results are reminiscent of key findings by Athey and Bagwell (2008) who study optimal collusion in repeated oligopolies in which firms incur serially correlated cost shocks.⁶ They show that under perfect persistence and 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 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, full compression holds regardless of the shape of the hazard rates and, more generally, for fairly arbitrary martingales.

To understand why compression holds in the global commons problem, the comparison with a collusive oligopoly or a procurement auction is instructive. As in these environments, incentive constraints are most restrictive for "low cost" types. In the present model, these are types with low consumption value (or high conservation value) for carbon. Information rents are therefore assigned to these types to prevent mimicry of higher cost types. When the participants are firms, the information rent is in the form of an expected market share. Incentive-constrained efficiency entails that higher production quotas are assigned to low cost types of seller. In this case, optimal mechanisms will not generally be compressed because allocation of market share inherently requires hard trade-offs between market shares of different types of firms. In the collusive oligopoly, for instance, a higher production quota assigned to one type of seller must be offset with a lower one to another.

By contrast, in the global commons problem there are no such trade offs. A country's "market share" is its expected net present value of "stored resource" which behaves like a public good. 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. Thus, a planner can increase each country's value of stored carbon simultaneously by reducing everyone's quota. Moreover, the free rider problem implies that any candidate for an optimal quota will typically push consumption below countries' individual incentives for extraction. Hence, if such a quota were non-compressed then the individual incentives for higher consumption (lower storage) induce low types to mimic higher types.

We characterize the compressed quota, and further show that it can be imple-

⁶See the subsequent Literature Review section for related results.

mented 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. Hence, compression on and off-path is part of a self enforcing agreement.

Finally, we discuss the role of side payments. In general side payments can mitigate inefficiency of full compression if countries that end up with low usage value are subsidized by those with high usage value. Because this may require that poor countries subsidize rich ones, however, we argue that efficiency-improving side payments may be infeasible.

The rest of the paper is organized as follows. The Section 2 summarizes the literature on dynamic mechanisms design as it applies to global commons. Section 3 describes the benchmark model of full information. In that model there are no shocks and each country's resource type is common knowledge. Section 4 introduces private persistent shocks. We first take up the case of perfect persistence, and then extend the results to the imperfect persistence case. Section 5 concludes with a discussion of potential implications for policy. Section 6 comprises an Appendix with the proofs.

2 Related Literature

The present paper draws upon two distinct literatures: the theoretical literature on dynamic mechanisms and the practical literature on regulating the commons and the environment. Both literatures are too vast to describe here and so we focus instead on narrow issues involving the intersection of the two.

Understandably much of the practical literature focuses on carbon regulation. Bodansky (2004) summarizes many of the hurdles faced by proposals to regulate emissions. A large quantitative literature has emerged to evaluate these. Some key quantitative assessments of carbon tax policies, for instance, include Nordhaus (2006, 2007) and Stern (2006), Golosov et al (2011), and Acemoglu et al (2012). Krusell and Smith (2009) calibrate a model of the global economy with fossil fuel use. They provide quantitative assessments of carbon taxation and cap and trade policies with the goal of achieving a zero emissions target. Rouillon (2010) proposes a competitive pricing scheme with directed transfers between individuals in a general common pool problem.

Most of these proposals focus on pricing mechanisms designed for firms and consumers in a competitive resource market. The present study, however, is concerned primarily with incentives of "large players", i.e., countries with strategic influence. In that vein, Barrett (2003) and Finus (2001) argue that any international mechanism

applied to countries must be dynamically self-enforcing. Consequently, they propose repeated game models in which international climate agreements are implementable in subgame perfect equilibria.

A number of models in the literature extend the self-enforcement constraint to non stationary commons games that better characterize the dynamics of resource use.⁷ Cave (1987) examines the traditional full information common pool model of Levhari and Mirman (1980). He shows that punishment strategies that trigger the Markov Perfect equilibrium (the “business-and-usual” outcome) can enforce full cooperation of the agreement when the participants are sufficiently patient. Dutta and Radner (2004, 2006, 2009) study a dynamic game model of energy usage with an emissions externality. Because the emissions stock is cumulative, dynamic incentives for energy usage may change over time.⁸ Dutta and Radner characterize the optimal emissions quota and, like Cave, compare it to the “business-as-usual” Markov Perfect benchmark. Battaglini and Harstad (2012) examine endogenous coalition formation in environmental agreements. In their model, global pollution can be addressed by investment in green technologies. The problem is that if agreement is “contractually complete” then countries may refuse to participate. Whereas if the agreement is incomplete, then it gives rise to an international hold up problem which, fortunately, can be mitigated when large coalitions of countries sign on to the appropriately structured agreement.

Like these, we stress the importance of ensuring that compliance be dynamically self-enforcing. The main difference is that in the aforementioned literature, there is no asymmetric information.

With the addition of informational asymmetries here, dynamic self-enforcement must apply to truthful disclosure as well compliance in any regulatory mechanism. In non-commons environments, there are such models and our study builds on them to an extent. For instance, Athey, Bagwell, and Sanchirico (2001), Aoyagi (2003), and Skyrpazc and Hopenhayn (2004) model optimal collusion of firms with private, iid shocks. Compression-like results hold in these models when hazard rates are monotone.⁹ The fact that compression holds here without the monotone hazard assumption reflects a basic difference between the commons problem and collusion/auction environments. We discuss these differences at length in Sect. 4.4. Studies of dynamic

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

⁸Strictly speaking, there is no common pool problem in their model since the marginal cost of energy usage is constant. Introducing convexity into the marginal extraction cost would bring their model closer to a global commons problem as modeled in the present paper.

⁹For related results, see McAfee and McMillan (1983) who model static collusion environments, and Amador, Angeletos, and Werning (2003) who model consumption-saving environments with iid shocks.

mechanisms with persistent shocks include Athey and Bagwell (2008) (discussed earlier), Pavan, A., I. Segal, and J. Toikka (2012), and Halac and Yared (2012). The latter demonstrate optimal rules that are explicitly uncompressed and, in fact, history dependent. Their setting - a model of dynamically inconsistent government - is quite different from either collusion or commons environments.

3 The Full Information Benchmark Model

3.1 Basic Setup

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

To better motivate the framework, we use carbon usage as a leading, albeit imperfect, example.¹⁰ The current stock of the resource at date t is given by ω_t . In the case of carbon, the current stock ω_t is the amount of “stored” greenhouse gas — the amount of carbon currently preserved under ground or in forest cover. Initially, we assume that the stock is known, and each country is able to precisely control its internal resource usage. We interpret the stock as a “sustainability” bound rather than an absolute quantity available. If the stock were to fall below its sustainability bound, the consequences would be disastrous. Fix the initial stock at $\omega_0 > 1$.

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

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

When $\gamma \leq 1$ the resource depreciates exponentially at rate γ . However, $\gamma > 1$ allows for growth in the stock. A transversality condition $\delta\gamma < 1$ is assumed to hold.

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

Let $\mathbf{c}_t = (c_{1t} \dots, c_{nt})$ denote the date t profile of resource consumption. The entire dynamic path profile of resource consumption is the given by

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

A path \mathbf{c} is then *feasible* if it is consistent with the technological constraint (1) and $C_t \leq \omega_t$ at each date t . Given a feasible consumption path \mathbf{c} , the long run payoff to country i is

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

Flow payoffs are discounted by δ each period. The value θ_i is the weight given to country i 's log consumption, whereas $1 - \theta_i$ is the weight assigned to the remaining resource stock $\omega_t - C_t$. The parameter θ_i is country i 's “resource type” or simply its “type” and is assumed to lie in an interval $[\underline{\theta}, \bar{\theta}] \subset [0, 1]$. We assume $n > \bar{\theta}/\underline{\theta}$, a condition which limits the global dispersion of types relative to the number of countries. A *type profile* is given by $\theta = (\theta_1, \dots, \theta_n)$. Following standard convention, $\theta_{-i} = (\theta_j)_{j \neq i}$.

Interpreting payoffs. 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 payoff weights both resource consumption and resource conservation. 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 “pure preference interpretation” builds on, and may be compared to, traditional “fish war” models of common pool resource usage dating back to Levhari and Mirman (1980). Those models assume $\theta_i = 1$ for all i , in which case a user of the resource merely trades off the value of 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.¹¹ In the case of carbon-based resources, countries obviously value the use of fossil fuels and timber, but recognize the associated GHG emissions as a costly by-product. Both benefits

¹¹The 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.

and costs of usage differ across countries. Warmer average temperatures resulting from GHG emissions are viewed differently in Greenland than in Sub-saharan Africa.

A second interpretation is that θ_i reflects production 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 the 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 common resource. The carbon input gets used up in the production process, while the non-carbon input is renewable (e.g., the “eco-system”) but depreciates/appreciates at rate γ according to (1). Each country utilizes the inputs at different intensities. Countries with larger θ_i use more of the carbon input to produce a given output. Richer countries, for instance, have larger carbon 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.¹²

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 . For now, there are no side payments between countries. We discuss the role of side payments in Section 4.5. Additionally, 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.

3.2 Optimal Quota Systems

Our interest is in international agreements chosen by the participants through a coordinating body such as the U.N. We refer to this body as the International Agency

¹²Size differentials would be captured instead by differential welfare weighting in any planner’s problem.

(IA). The IA, as envisioned here, operates by the consent of its members, gathers and makes available information, makes recommendations, and suggests sanctions for violations. It does not have the ability to enforce sanctions.

Instead, the international agency's recommendation takes the form of a *quota system*. By a quota system, we mean only a mapping $\mathbf{c}^*(\theta)$ from type profiles to feasible consumption paths. Hence, $c_{it}^*(\theta)$ is the targeted consumption recommended for country i at date t given the global type profile θ .

Because the IA cannot directly impose or enforce a quota system, its recommended quota system must be *implementable* by a subgame perfect equilibrium (SPE) of the dynamic resource game. In order to define the implementability requirement, a few extra bits of notation are needed. Let $h^t = (\omega_0, \mathbf{c}_0, \omega_1, \mathbf{c}_1, \dots, \omega_{t-1}, \mathbf{c}_{t-1}, \omega_t)$ be the date t history of usage levels and resource stocks, including the current stock ω_t . The initial history is $h^0 = \omega_0$. A *usage strategy* $\sigma_i(h^t, \theta)$ for country i maps histories and global type profiles to desired resource consumption c_{it} at date t . A *usage profile* is given by $\sigma = (\sigma_1, \dots, \sigma_n)$. A default rule guarantees that usage profiles generate feasible consumption paths.¹³

The long run payoff of usage profile σ to country i after any history h^t is expressed recursively as

$$V_i(h^t, \sigma, \theta) = \theta_i \log \sigma_i(h^t, \theta) + (1 - \theta_i) \log \left(\omega_t - \sum_{j=1}^n \sigma_j(h^t, \theta) \right) + \delta V_i(h^{t+1}, \sigma, \theta), \quad (3)$$

where decisions in t determine the history h^{t+1} entering $t + 1$.¹⁴ A profile σ is a *subgame perfect equilibrium (SPE)* if each σ_i maximizes country i 's long run payoff V_i after every history h^t .

An SPE profile σ may then be said to *implement* a quota system \mathbf{c}^* if σ generates consumption \mathbf{c}^* along the equilibrium path.¹⁵ Hence, if a SPE profile σ implements a quota system $\mathbf{c}^*(\theta)$, then it follows from (2) and (3) that $V_i(h^0, \sigma, \theta) = U(\omega_0, \mathbf{c}^*(\theta), \theta_i)$.

The IA's problem can now be formally stated as one that recommends a quota $\mathbf{c}^*(\theta)$ and a subgame perfect profile σ that solves

$$\max_{\mathbf{c}^*(\theta)} \sum_{i=1}^n U_i(\omega_0, \mathbf{c}^*(\theta), \theta) \quad \text{such that } \mathbf{c}^*(\theta) \text{ is implemented by the SPE } \sigma. \quad (4)$$

¹³For our purposes, any number of default rules will work. For instance, let $c_{it} = \sigma_i(h^t, \theta)$ if $\sum_j \sigma_j(h^t, \theta) < \omega_t$, and let $c_{it} = \frac{\sigma_i(h^t, \theta)}{\sum_j \sigma_j(h^t, \theta)} \omega_t$ otherwise.

¹⁴Formally, $h^{t+1} = (h^t, \sigma(h^t, \theta), (\omega_t - \sum_{j=1}^n \sigma_j(h^t, \theta))^\gamma)$.

¹⁵Formally, \mathbf{c}^* is generated recursively: $\mathbf{c}_0^*(\theta) = \sigma(h^0, \theta)$, then $\mathbf{c}_1^*(\theta) = \sigma(h^0, \mathbf{c}_0^*(\theta), (\omega_0 - C_0^*(\theta))^\gamma, \theta)$, and so forth...

Hence, 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 (4) 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 (4) can be found by breaking the problem into two steps.

- Step 1 characterizes the optimal quota without the compliance constraints. In the full information case, this amounts to solving (4) without the equilibrium constraint. This “relaxed” problem yields the unconstrained, “first-best” quota chosen by an international agency with the ability to impose and enforce its choice upon the participants.
- Step 2 shows that the first-best quota can be implemented by SPE profile σ^* .

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

To characterize the first-best quota (Step 1 in the above algorithm), it turns out to be easier to work with extraction rates rather than levels. For any dynamic consumption path \mathbf{c} , let \mathbf{e} denote the corresponding extraction rates, defined by $e_{it} = \frac{c_{it}}{\omega_t}$. 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. The first order condition for the IA's relaxed problem in e_{it} is:

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

(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 combining it with the first order condition yields the IA's Euler equation

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




Figure 1: First Best Extraction Rates

The forward solution to the IA's Euler equation is then easily found to be

$$\frac{\theta_i(1 - \mathcal{E}_t)}{e_{it}} - \sum_j (1 - \theta_j) = \frac{n\delta\gamma}{1 - \delta\gamma} \quad (5)$$

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 first-best aggregate extraction rate is a fraction $(1 - \delta\gamma)$ of the average resource type $\frac{\Theta}{n}$. This aggregate rate is achieved by country-specific extraction rates (see Fig. 1).

$$e_i^*(\theta) = \frac{\theta_i(1 - \delta\gamma)}{n}. \quad (6)$$

These rates yields a quota system \mathbf{c}^* given by

$$c_{it}^*(\theta) = \frac{\theta_i(1 - \delta\gamma)}{n} \omega_0^{\gamma t} \left(1 - \frac{\Theta(1 - \delta\gamma)}{n}\right)^{\gamma(1 - \gamma^t)/(1 - \gamma)} \quad (7)$$

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 in (7) are increasing in one's own resource type θ_i but decreasing in the cross-country average Θ/n . Hence, countries

with larger-than-average carbon usage intensity can extract more. Put another way, “pro-consumption” types should extract more while “pro-conservation” types should extract less.

Compliance with the first best quota is not automatic. One could compare the IA’s Euler equation to that of each individual country. That latter corresponds to a country’s individual incentives for resource consumption, and is given by:¹⁶

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

As before, the forward solution is easily found.

$$\frac{(1 - \mathcal{E}_t)}{e_{it}} = \frac{\delta\gamma}{\theta_i(1 - \delta\gamma)} + \frac{1 - \theta_i}{\theta_i} \quad (8)$$

From (8), one can calculate country’s best response $e_{it} = BR_i(\theta_i, \mathcal{E}_{-it})$ to the extraction rates \mathcal{E}_{-it} of other countries:

$$BR_i(\theta_i, \mathcal{E}_{-it}) = \theta_i(1 - \delta\gamma)(1 - \mathcal{E}_{-it}) \quad (9)$$

Predictably, pro-extraction types have a greater incentive to extract more. It is easy to verify that $BR_i(\theta_i, \mathcal{E}_{-i\ t}^*) > e_i^*$. In other words, the individual country’s incentives are toward greater extraction than that prescribed by the IA.¹⁷

Consequently, the quota system described by (7) is the optimal one only if it can be implemented by a SPE (Step 2 in the solution algorithm). This is summarized in the following lemma.

¹⁶Starting, as before, with each country’s first order condition

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

then differentiate the value function for country i , giving

$$\frac{\partial U_i}{\partial \omega_t} = \frac{1}{\omega_t} + \delta\gamma \frac{\partial U_i}{\partial \omega_{t+1}} \omega_t^{\gamma-1} (1 - \mathcal{E}_t)^\gamma.$$

Iterating this second equation forward one period, then plugging in the FOC to substitute out $\frac{\partial U_i}{\partial \omega_{t+2}}$ on right-hand side (iterated), gives an expression for $\frac{\partial U_i}{\partial \omega_{t+1}}$. Substituting back into original FOC yields the Euler equation above.

¹⁷Note that solving (9) as a simultaneous system of equations across countries yields the unique Markov Perfect equilibrium (MPE) — the analogue of the “business-as-usual equilibrium” computed by Cave and by Dutta and Radner in their models. The analytical solution to the MPE in the present model is non-trivial. A full characterization is worked out in Harrison and Lagunoff (2013).

Proposition 1 (Full Information Benchmark) *For any profile θ of resource types, the IA’s optimal quota system (the solution to (4)) is the quota system $\mathbf{c}^*(\theta)$ described by (7).*

Proposition 1 is useful benchmark for later the results with private information. Note that the implementation holds for *any* discount factor $\delta > 0$, including impatient ones. This is possible because payoffs are unbounded below, reflecting the idea that in the case of global commons, the costs of resource depletion may be catastrophic. As a consequence, the IA can recommend further threats of resource depletion in any continuation payoff, even ones that are already punitive, to enforce compliance. Since increased resource depletion hurts all countries, the credibility of the punishment depends on even harsher punishment if the countries fail to carry out the sanction. This means, in turn, that increasingly severe depletion threats must be used, each such threat made credible only by even more severe depletion threats later on, and so forth. The successive threats are only reached, of course, by further deviations at each counterfactual stage. Since each country’s payoff in the residual stock is unbounded, the sequence of threats can be recursively defined. The proof in the Appendix gives the formal details.

Note further 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.

4 Persistent Private Shocks

Our main results pertain to the case where a country’s internal costs and benefits of resource usage are privately observed and change over time as countries are hit with serially correlated, idiosyncratic shocks.

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

To sketch out one example, consider a scenario in which a warmer climate lowers a country’s agricultural yields. This, in turn, leads to more intense use of petroleum-based fertilizer, thus increasing its relative value of carbon. In another scenario, the warmer climate leads to an increase in the saline contamination of the country’s fresh

water fisheries. In that case, the relative value of carbon decreases as GHG costs increase. The international agency is keen to know which scenario prevails, but must rely on self-reported information by the countries.¹⁸

4.1 The Shock Structure

To capture these issues in the simplest possible way, we allow the resource type θ to vary over time. Let θ_{it} denote country i 's resource type at date t . The distribution over θ_{it} is given by a Markov kernel, $F_i(\theta_{it}|\theta_{i,t-1})$, given the country's prior resource type, $\theta_{i,t-1}$. The distribution $F_i(\cdot|\theta_i)$ is differentiable in θ_i , has full support on $[\underline{\theta}, \bar{\theta}]$, and admits a continuous, conditional density $f_i(\cdot|\theta_i)$ for each $\theta_i \in [\underline{\theta}, \bar{\theta}]$. Denote i 's initial type distribution by $F_{i0}(\theta_{i0})$, and i 's marginal distribution (at date 0) on θ_{it} by $F_{it}(\theta_{it})$.

The shocks are IID across countries, and while the Markov distribution F_i for country i is commonly known to all countries, each country's realized shock each period is privately observed. The profile of types in date t is

$$\theta_t = (\theta_{1t}, \theta_{2t}, \dots, \theta_{it}, \dots, \theta_{nt}).$$

Its conditional distribution is given by $F(\cdot|\theta_{t-1})$. The joint conditional distribution excluding country i is given by $F_{-i}(\cdot|\theta_{-i,t-1})$. Finally, let $\theta^t = (\theta_0, \theta_1, \dots, \theta_t)$ be the history of realized profiles up to and including date 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 given by the sequence $\mathbf{c}^* = \{c_t^*(\theta^t)\}$, $t = 0, 1, \dots$. The IA's recommendation at each date now depends on the entire history of shocks θ^t up to that point.¹⁹

To make an effective recommendation, the international agency is a gatherer and dispenser of information. At the beginning of each period t , the IA solicits information concerning each country's realized type θ_{it} . As with resource consumption, the IA cannot enforce truthful disclosure by its members. Instead, the IA serves as a vehicle

¹⁸This is, in fact, the case. See Article 12 under the UNFCCC, and Articles 5 and 7 under the Kyoto Protocol.

¹⁹To see why the type history, observe that the initial recommendation c_0^* may vary with the initial realization θ_0 . Next period's recommendation c_1^* varies, of course, with θ_1 . To be feasible, however, c_1^* must also depend on the stock ω_1 which is determined, in turn, by c_0^* . Rolling things back from an arbitrary date t , the current recommendation may vary with the type history θ^t .

for coordinating information and usage. Each member country chooses whether or not to disclose its type (as, for instance, when countries make public their national income accounts, estimates, and forecasts). A country's reported type is denoted by $\tilde{\theta}_{it}$. The entire profile of types reported by all countries is $\tilde{\theta}_t = (\tilde{\theta}_{1t}, \dots, \tilde{\theta}_{nt})$.

Though not intended as a model of existing international agreements, the disclosure mechanism nevertheless resembles some features of existing mechanisms. Article 12 of the UN Framework Convention on Climate Change, for instance, requires its signatories to periodically submit, among other things, a "national inventory of anthropogenic emissions," and a "specific estimate of the effects that the policies and measures ... will have on anthropogenic emissions by its sources and removals by its sinks of greenhouse gases..."

In order to obtain the desired quota system, the IA suggests a strategy profile that now includes a disclosure option. Formally, let $\tilde{\theta}^t = (\tilde{\theta}_0, \tilde{\theta}_1, \dots, \tilde{\theta}_t)$ denote the *disclosure history* up to date t . A *disclosure strategy* for country i is map $\mu_i(h^t, \tilde{\theta}^{t-1}, \theta_{it}) = \tilde{\theta}_{it}$ describing i 's report 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 now a 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 $\mu = (\mu_1, \dots, \mu_n)$ and $\sigma = (\sigma_1, \dots, \sigma_n)$ 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

$$V_i(h^t, \tilde{\theta}^t, \sigma, \mu | \theta_{it}) \equiv \int_{\theta_{-it}} \left[\theta_{it} \log \sigma_i(h^t, \tilde{\theta}^t, \theta_{it}) + (1 - \theta_{it}) \log(\omega_t - \sum_{j=1}^n \sigma_j(h^t, \tilde{\theta}^t, \theta_{jt})) \right. \\ \left. + \delta \int_{\theta_{it+1}} V_i(h^{t+1}, \tilde{\theta}^{t+1}, \sigma, \mu | \theta_{it+1}) dF_i(\theta_{it+1} | \theta_{it}) \right] dF_{-i}^*(\theta_{-it} | h^t, \tilde{\theta}^t) \quad (10)$$

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 (μ, σ) yields \mathbf{c}^* along the outcome path.²⁰

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 $F_i^*(\theta_{it} | h^t, \tilde{\theta}^t)$, $i = 1, \dots, n$ constitute a *Perfect Bayesian equilibrium (PBE)* if (i) at the consumption stage in date t , σ_i and μ_i together maximize i 's long run expected payoff (defined in (10)) 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.

Note that at the disclosure stage, countries contemplate disclosure deviations from the prescribed plan, taking account of the fact that they have the freedom to deviate 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.

If, however, (σ, μ) is a truth-telling PBE that implements \mathbf{c}^* , then along the equilibrium path, we have

$$V_i(h^t, \theta^t, \sigma, \mu | \theta_{it}) = U_i(\omega_t, c^*(\theta^t), \theta_{it}) \equiv \theta_{it} \log c_{it}^*(\theta^t) + (1 - \theta_{it}) \log(\omega_t - C_t^*(\theta^t)) + \delta \int_{\theta_{t+1}} U_i(\omega_{t+1}, c^*(\theta^{t+1}), \theta_{it+1}) dF(\theta_{t+1} | \theta_t) \quad (11)$$

which is country i 's *realized* payoff under quota system \mathbf{c}^* after any history θ^t . Country i 's *ex ante* expected payoff is therefore $\int_{\theta_0} U_i(\omega_0, c^*(\theta_0), \theta_{i0}) dF_0(\theta_0)$.

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

²⁰As with the full information case, the consumption path is generated recursively: $\mathbf{c}_0^*(\theta^0) = \sigma(h^0, \theta^0, \theta_0)$, then $\mathbf{c}_1^*(\theta^1) = \sigma(h^1, \theta^0, \theta_1)$, where $h^1 = (h^0, \sigma(h^0, \theta_0, \theta_0), (\omega_0 - C_0(\theta^0))^\gamma)$, etc., ...

solves

$$\max_{\mathbf{c}^*} \sum_i \int_{\theta_0} U_i(\omega_0, c^*(\theta_0), \theta_{i0}) dF_0(\theta_0) \quad \text{subject to } \mathbf{c}^* \text{ implemented by the PBE } (\sigma, \mu). \quad (12)$$

Any quota that solves (12) will be called an *optimal quota*. Our main result is that the optimal quota has a special form which we refer to as *fully compressed*. A quota system \mathbf{c}^* 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 (12) into two parts, much as we did in the full information model.

- Step 1, as before, characterizes the solution to a “relaxed problem” without the compliance constraints. We characterize the closed-form solution to Step 1, and show that it is, in fact, a fully compressed quota.
- Step 2 shows that the fully compressed “relaxed 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 or off the equilibrium path.

4.3 Perfect Persistence

For illustrative purposes, we first take 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 , and drop all time subscripts on the shock distributions. This is a reasonable approximation to an situation where environmental change is more rapid than technological progress. We return to the general imperfect persistence model in Section 4.4.

Proposition 2 *Consider the model with private, perfectly persistent shocks. If the effective discount factor $\delta\gamma$ is sufficiently close to one, then the optimal quota \mathbf{c}^* (the solution to (12)) is fully compressed. Specifically,*

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

$$e_i^* = \frac{(1 - \delta\gamma) \left[\int_{\underline{\theta}}^{\bar{\theta}} \theta_i dF_i(\theta_i) \right]}{n} \quad (13)$$

for each country i .

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

The Proposition describes the optimal quota by its extraction rates rather than by its consumption levels. For completeness, the optimal consumption is given by

$$c_{it}^*(\theta^t) = e_i^* \omega_t^* = \frac{(1 - \delta\gamma) \int_{\underline{\theta}}^{\bar{\theta}} \theta_i dF_i(\theta_i)}{n} \omega_0^{\gamma^t} \left(1 - \frac{(1 - \delta\gamma) \sum_i \int_{\underline{\theta}}^{\bar{\theta}} \theta_i dF_i(\theta_i)}{n} \right)^{\gamma(1-\gamma^t)/(1-\gamma)} \quad (14)$$

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

To illustrate what compression means for the optimal quota, compare the optimal rate \mathbf{e}^* in (13) to the full information solution in (6). Take a particularly simple case, suppose that $F_i = F_j = F$, i.e., all countries draw their shocks from the same initial distribution. Then the optimal quota under private shocks assigns all countries *identical* extraction rates, despite the fact that the international agency can condition its recommendation on the information disclosed by each of the countries. This is illustrated in Figure 2. If, for instance, F is uniform on $[\underline{\theta}, \bar{\theta}]$, then in the optimal quota, all countries are required to extract resources at the same rate $e_i^* = (1 - \delta\gamma)(\bar{\theta} + \underline{\theta})/2n$, even if they realize very different resource types.

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

4.4 The Basic Logic

Here, we outline the basic logic of the argument. The next section discusses some extensions and robustness checks of the result. A detailed proof of Proposition 2 is in the Appendix.




Figure 2: Optimal quota: Private Shocks vs Full Information

Part 1 of the Proof solves a “relaxed problem” similar to the one in the full information benchmark. In this case, the PBE constraint in the relaxed problem is replaced by a weaker constraint that requires only truthful disclosure. Compliance incentives following disclosure are ignored. Part 2 of the Proof then shows that the relaxed-optimal quota can be implemented by a full-blown PBE in which a country’s prescribed usage rate after any history is stationary and compressed. For now, we focus on the intuition for Part 1.

In the relaxed problem, the International Agency faces only a single disclosure stage at the beginning of $t = 0$ (since compliance is ignored). This is given by the truth-telling constraint:

$$\int_{\theta_{-i}} U_i(\omega_0, \mathbf{c}^*(\theta), \theta_i) dF_{-i}(\theta_{-i}) \geq \int_{\theta_{-i}} U_i(\omega_0, \mathbf{c}^*(\tilde{\theta}_i, \theta_{-i}), \theta_i) dF_{-i}(\theta_{-i}) \quad \forall \tilde{\theta}_i \quad \forall i \quad (15)$$

To understand logic of the relaxed problem, a comparison with auctions is instructive. Using the resource technology and some algebra, a country’s interim payoff under truth-telling (the left-hand side of (15)) can be expressed as²¹

$$\int_{\theta_{-i}} U_i(\omega_0, \mathbf{c}^*(\theta), \theta_i) dF_{-i}(\theta_{-i}) = \int_{\theta_{-i}} [R_i(\theta_i) - \theta_i Q_i(\theta_i)] dF_{-i}(\theta_{-i}) \quad (16)$$

²¹The formal derivation is found in the proof of Proposition 2 in the Appendix.

where

$$R_i(\theta_i) \equiv \frac{\log \omega_0}{1 - \delta\gamma} + \int_{\theta_{-i}} \sum_{t=0}^{\infty} \delta^t \sum_{j=0}^t \gamma^{t-j} \log(1 - \mathcal{E}_j^*(\theta)) dF_{-i}(\theta_{-i}) \quad \text{and} \quad (17)$$

$$Q_i(\theta_i) \equiv \sum_{t=0}^{\infty} \delta^t \log \left(\frac{1 - \mathcal{E}_t^*(\theta)}{e_{it}^*(\theta)} \right). \quad (18)$$

The payoff has a simple interpretation. $R_i(\theta_i)$ is the country's long run expected benefit of resource conservation, whereas $\theta_i Q_i(\theta_i)$ is its long run cost. 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 right-hand side of (16) clearly resembles as the interim expected profit of a firm in an oligopolistic industry or a seller in a procurement auction. Using the oligopolistic industry as a guide, $R_i(\theta_i)$ would correspond to an expected revenue/compensation, θ_i would be i 's marginal cost, and $Q_i(\theta_i)$ its market share. Standard results from optimal auctions (Myerson (1981)) may be applied to show that the expected payoff to this firm/country in an incentive compatible quota system (one satisfying (15)) is

$$R_i(\bar{\theta}) - \bar{\theta} Q_i(\bar{\theta}) + \int_{\underline{\theta}}^{\bar{\theta}} \frac{F_i(\theta_i)}{f_i(\theta_i)} Q_i(\theta_i) f_i(\theta_i) d\theta_i \quad (19)$$

where Q_i must be weakly decreasing in the type θ_i .

In other words, a country's expected payoff in an incentive compatible quota system is its payoff as the highest resource type $\bar{\theta}$, plus the usual "information rent" required to induce truth-telling. The information rent is increasing in a weighted average of the "market share" term Q_i . Critically, in most seller environments there are trade offs in the production quota $Q_i(\theta_i)$ between sellers 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.²²

One exception to this is when the hazard rate F_i/f_i is strictly increasing. In that case, the loss from inefficient production is offset by allocative concerns of the planner. Athey and Bagwell (2008) for instance, analyze a repeated oligopoly setting in which firms receive serially correlated cost shocks each period. They show that

²²See Pavan, Segal, and Toikka (2012) for a comprehensive characterization of optimal mechanisms in Markov models of private information.

the optimal production quota for colluding firms is fully compressed (which they refer to as “rigid”) when either hazard rates are increasing or the maximum possible compensation from monopoly pricing is large enough.²³ McAfee and McMillan (1992) show a similar result when buyers collude in a static procurement auction.

In the present model, full compression holds even though hazard rates need not be monotone. The difference between these results and our own concerns the nature of Q_i . Recall from (18) that $Q_i(\theta_i)$ country i ’s expected present value of stored resource relative to its own consumption. Because the stored resource is non-rivalrous, there are no tangible trade offs. All countries’ values for, say, carbon storage can be increased simultaneously by simply withholding carbon consumption.

Indeed, this is essential because the free rider problem will lead to under-provision of storage — or over-provision of consumption. In particular, it can be shown that any quota that maximizes the sum of incentive-adjusted payoffs in (19) must assign stationary extraction rates below the full information-optimal extraction rate, $\bar{\theta}(1 - \delta\gamma)/n$, for the highest type $\bar{\theta}$. Given this bound, when countries are sufficiently patient then any candidate for an incentive-constrained optimal quota will typically push consumption below each countries’ individual incentives for over-extraction. This tension between individual incentives and social optimum is illustrated in Fig. 3. If any candidate for an optimal quota were non-compressed then the individual incentives for higher consumption (lower storage) would induce low types to mimic higher types.

Part 2 of the Proposition concerns the implementation of the relaxed solution. To implement the quota \mathbf{e}^* one might think that the IA could simply use the sanctions constructed in the full information case. That logic (outlined in the proof of Proposition 1) requires credible punishment at each counterfactual stage depends on even harsher punishment if the countries fail to carry out the sanction.

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

²³Roughly, a rigid/compressed production quota can be shown to stochastically dominate any strictly decreasing one when weighted by an increasing hazard rate. Their collusive mechanism allocates shares of production to firms who service an inelastic demand. The production shares must, on average, add to one.

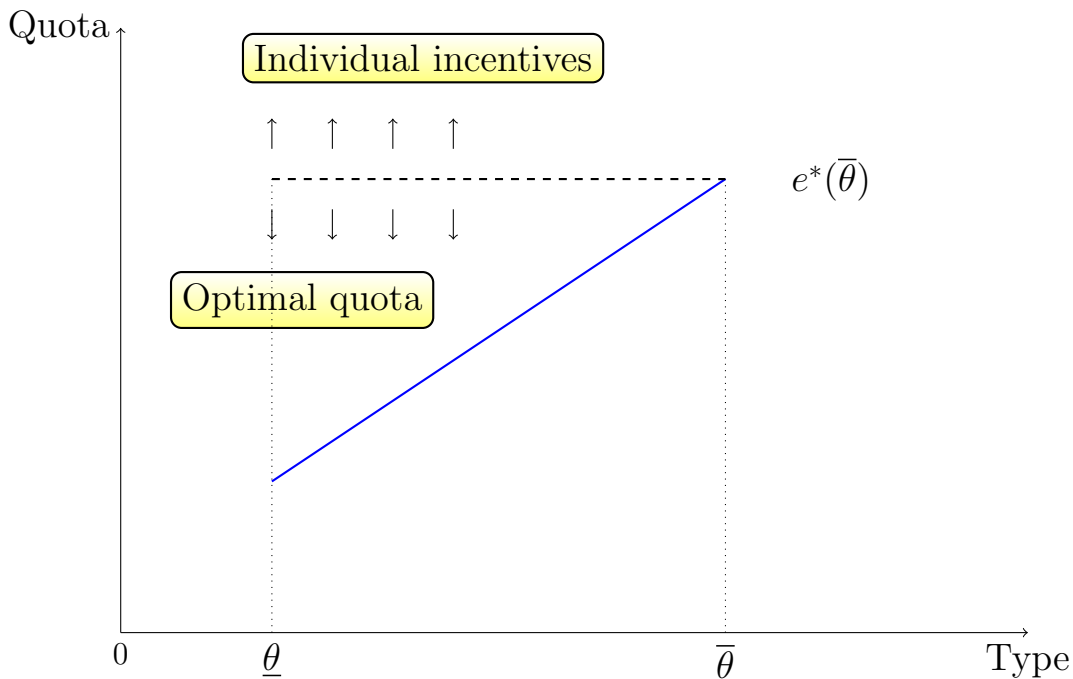


Figure 3: Individual incentives vs IA’s “relaxed” optimum

the Appendix for details).

4.5 Robustness Issues: commitment and side payments

The basic logic above may be applied to the general model of imperfect persistence as well. In light of the logic, two extensions are worth considering before moving on to the general model.

Commitment. We have supposed the existence of a “U.N.-like” international agency that can fully commit to a mechanism at $t = 0$. At first glance, this may seem critical. Consider, for instance, a truth-telling equilibrium that implements the optimal quota. Clearly, since the quota is fully compressed, each country has no incentive to lie about its type in date $t = 0$. This means that the IA now has full information about countries’ types heading into date $t = 1$. A dynamically consistent IA would make use of this information to implement the first-best solution in date $t = 1$ and beyond. However, this destroys the initial incentive for truthful disclosure.²⁴

²⁴This “ratchet effect” was originally observed by Roberts (1984) in a dynamic Mirleesian model of optimal taxation.

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

Side payments. The model currently does not allow for transfers or side payments between countries. The literature on self-enforcing environmental agreements (absent private information) includes a fair number of models (e.g., Finus (2001)) that allow for monetary transfers or side payments.²⁵

On the one hand, the absence of side payments in the present analysis limits its applicability since there are well established results dating from d’Aspremont and Gerard-Varet (1979) that show that balanced transfer schemes exist that can implement the first best allocations in environments with private information. Formally, let $s_i(\tilde{\theta})$ denote a tax/transfer from/to country i . Here, $s_i(\tilde{\theta}) > (<)0$ connotes a monetary transfer to (tax paid by) country i given reported profile $\tilde{\theta}$. A *balanced* transfer scheme satisfies $\sum_i s_i(\theta) = 0$ for every profile θ .

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

Note that transfers do not affect incentives at the compliance stage. Hence, the proof in the Appendix constructs a transfer scheme that adapts the static d’Aspremont and Gerard-Varet scheme to the present model.

Despite the apparent reversal of the results suggested in the Proposition, we argue that side payment schemes are not likely to be politically feasible in any global commons arrangement. They face two main problems. First, implementations with side payments such as the one above typically require no budget constraints.

This is all the more serious in the present model since any welfare improving transfer scheme will at some point likely involve transfers from poor countries to rich ones. Specifically, in order to counteract compression, countries that realize shocks with high usage value (relative to its value of conservation) would be required to subsidize those with low usage value. Yet, the large projected increase in usage by developing countries, coupled with per capita reductions in GHG production by wealthier countries, means that relatively poorer countries will eventually be the ones with high usage value.²⁶ It seems unlikely that any agreement would require poorer countries to make payments to their wealthier counterparts.

²⁵Side payments may include non-climate benefits such as technology transfers, trade terms, and so forth. Their inclusion in a climate agreement is often referred to as *issue linkage*.

²⁶See projected usage estimates from the Carbon Dioxide Information Analysis Center (CDIAC), cdiac.ornl.gov.

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

4.6 Imperfectly Persistent Private Shocks

We now return to the general model with imperfectly persistent shocks. The main result is:

Proposition 4 *Consider the model with private, imperfectly persistent shocks. Suppose that for each country i , F_i is assumed to satisfy*

$$\theta_{it} = \int \theta_{it+1} dF_i(\theta_{it+1} | \theta_{it}). \quad (20)$$

If the effective discount factor $\delta\gamma$ is sufficiently close to one, then the optimal quota \mathbf{c}^ (the solution to (12)) is fully compressed. Specifically,*

- (1) *the extraction rates corresponding to \mathbf{c}^* are fully compressed and non-stationary. In particular, for each country i , and each date t ,*

$$e_{it}^* = \frac{(1 - \delta\gamma) \int_{\underline{\theta}}^{\bar{\theta}} \theta_{it} dF_{it}(\theta_{it})}{n} \quad (21)$$

for each country i , and

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

Equation (20) entails that the stochastic process $\{\theta_{it}\}$ is a martingale. We later refer to it as the “no-drift” assumption. Clearly, the end result coincides with the perfect persistence model when types are perfectly correlated across time. The basic logic of the first part of Proposition 2 is similar to the perfect persistence case. That

is, we first solve the following “super-relaxed” planner’s problem:

$$\begin{aligned}
& \max_{c^*} \sum_i \int_{\theta_0} U_i(\omega_0, c^*(\theta_0), \theta_{i0}) dF_0(\theta_0) \quad \text{subject to} \\
& \int_{\theta_{-it}} U_i(\omega_t, c^*(\theta^t), \theta_{it}) dF_{-i}(\theta_{-it} | \theta_{t-1}) \geq \int_{\theta_{-it}} U_i(\omega_t, c^*(\theta^t \setminus \tilde{\theta}_{it}), \theta_{it}) dF_{-i}(\theta_{-it} | \theta_{t-1}) \\
& \forall \theta^t \forall \tilde{\theta}_{it} \forall i \forall t
\end{aligned} \tag{22}$$

Here, $\theta^t \setminus \tilde{\theta}_{it}$ is the history of type-profiles in which country i inserts $\tilde{\theta}_{it}$ in place of θ_{it} . This planner’s problem is “super-relaxed” in the sense that it not only ignores the compliance constraints, it also ignores multi-period deviations in disclosure. Only one-period deviations in the disclosure decisions are considered. The proof in the Appendix shows that the solution to this “super-relaxed” problem is a fully compressed quota given by equation (21). A fully compressed quota is shown to be immune from multi period deviations as well.

Proposition (4) shows that full compression holds in cases of imperfect persistence as well. Two countries that start identically will typically evolve very different benefits and costs of resource use. They are, nevertheless, prescribed the same consumption quota throughout.

To prove this, we once again exploit a decomposition of payoffs. As with perfectly persistent shocks, there is a “compensation” term and a “market share” term. With imperfect persistence, however, there is also a third term. 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. The “no-drift” assumption allows one to apply recursive logic. Namely, at each date, this third term disappears whenever future values of the quota are chosen to be fully compressed. We then apply Bellman’s Principle to show that the current value is maximized by a compressed quota.

As with the perfect persistence case, Part 2 of the Proposition asserts the existence of an implementing PBE that uses compressed punishments. The proof largely mimics the steps of the proof of Proposition 2 (part 2), and so we omit the details.

5 Conclusion

This paper studies dynamic mechanisms for global commons with environmental externalities. Using carbon consumption as the leading example, we generalize the dynamic resource extraction game of Levhari and Mirman to allow for direct, heterogeneous benefits of resource conservation across countries. 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.

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

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

These results differ from related findings by Athey and Bagwell (2008) and others in models of optimal collusion in markets. In those models, compression depends largely on shapes of the shock distributions. These differences raise natural questions about how the nature of the good — private versus public — plays a critical role in the optimal design.

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

The results may be reconciled to these approaches to the extent that transfers can be used, or that some information about local shocks is globally observed. We argued earlier, however, that the use of transfers or side-payments for allocating carbon is limited by budget and political constraints, particularly given the projected pattern of carbon usage between countries.

As for information, inefficient compression could be mitigated if the data gathering process occurs above and beyond the reach of sovereign filters. This could be

presumed of the climate science itself. As before, any mitigating mechanism based on globally public information must be self enforcing at compliance stage or should be enforceable by the International agent. In our view, the disclosure of country-specific observations of economic costs and benefits remains problematic because these rely to a large extent on each country's yearly disclosure of its national income accounts. Public provision of "investigative" resources to generate such information, however, might overcome this problem. It seems an excellent topic for future study.

6 Appendix: Proofs of the Results

Proof of Proposition 1. Fix a profile θ . As we have already shown that $\mathbf{c}^*(\theta)$ maximizes $\sum_i U_i$ without the equilibrium constraint, it remains to show that $\mathbf{c}^*(\theta)$ can be implemented by a SPE.

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

$$\begin{aligned} & \frac{\omega_0}{1 - \delta\gamma} + \sum_{t=0}^{\infty} \delta^t \left[\left(\frac{1}{1 - \delta\gamma} - \theta_i \right) \log(1 - \mathcal{E}_t) + \theta_i \log e_{it} \right] \\ \equiv & \frac{\omega_0}{1 - \delta\gamma} + \sum_{t=0}^{\infty} \delta^t u_{it} \end{aligned} \quad (23)$$

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

$$V_i^*(\theta) \equiv \frac{\omega_0}{1 - \delta\gamma} + \frac{1}{1 - \delta} \left[\left(\frac{1}{1 - \delta\gamma} - \theta_i \right) \log(1 - \mathcal{E}^*(\theta)) + \theta_i \log e_i^*(\theta) \right] \quad (24)$$

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 $V_i^0(\theta) = V_i^*(\theta)$.

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 (23) of an arbitrary flow payoff, the payoff in state τ is given by

$$V_i^\tau(\theta) = \frac{1}{1 - \delta\rho} \left[\left(\frac{1}{1 - \delta\gamma} - \theta_i \right) \log(1 - \mathcal{E}^\tau(\theta)) + \theta_i \log e_i^\tau(\theta) \right] + \frac{\delta(1 - \rho)}{1 - \delta\rho} V_i^{\tau-1}(\theta) \quad (25)$$

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 (25), one can show that the incentive constraint holds if for all i ,

$$V_i^{\tau-1}(\theta) \geq \frac{1 - \delta\rho}{1 - \delta\rho - \delta^2(1 - \rho)} \max_{e_i} \left[\left(\frac{1}{1 - \delta\gamma} - \theta_i \right) \log(1 - \mathcal{E}_{-i}^{\tau-1}(\theta) + e_i) + \theta_i \log e_i \right] + \frac{\delta^2}{1 - \delta\rho - \delta^2(1 - \rho)} \left[\left(\frac{1}{1 - \delta\gamma} - \theta_i \right) \log(1 - \mathcal{E}^\tau(\theta)) + \theta_i \log e_i^\tau(\theta) \right] \quad (26)$$

Clearly, these inequalities (one for each country) can always be made to hold by choosing \mathcal{E}^τ 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 Part 1 of Proposition 2.

The proof proceeds by first solving a “relaxed planner’s problem” in which the PBE implementation constraint is replaced by a simple truth-telling constraint,

$$\int_{\theta_{-i}} U_i(\omega, c^*(\theta), \theta_i) dF_{-i} \geq \int_{\theta_{-i}} U_i(\omega, c^*(\tilde{\theta}_i, \theta_{-i}), \theta_i) dF_{-i} \quad \forall \theta_i \quad \forall \tilde{\theta}_i \quad (27)$$

The “relaxed planner’s problem” is stated as

$$\max_{\mathbf{c}^*} \sum_i \int_{\theta} U_i(\omega, c^*(\theta), \theta_i) dF(\theta) \quad \text{subject to (27)}. \quad (28)$$

Since the relaxed problem ignores off-path deviations, all the subsequent arguments here refer to the on-path incentive constraints in (27).

Let \mathbf{c}^* denote a candidate solution to (28) and let \mathbf{e}^* denote the rates corresponding to \mathbf{c}^* . Notice that if an equilibrium (σ, μ) which implements \mathbf{c}^* satisfies the truth-telling constraint (27), 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,

$$\begin{aligned}
U_i(\omega_0, \mathbf{c}^*(\tilde{\theta}), \theta_i) &= \sum_{t=0}^{\infty} \delta^t \left[\theta_i \log \omega_t e_{it}^*(\tilde{\theta}) + (1 - \theta_i) \log \omega_t (1 - \mathcal{E}_t^*(\tilde{\theta})) \right] \\
&= \sum_{t=0}^{\infty} \delta^t \left[\log \omega_t + \theta_i \log e_{it}^*(\tilde{\theta}) + (1 - \theta_i) \log(1 - \mathcal{E}_t^*(\tilde{\theta})) \right] \\
&= \sum_{t=0}^{\infty} \delta^t \log \left(\omega_0^{\gamma^t} \prod_{j=0}^{t-1} (1 - \mathcal{E}_j^*(\tilde{\theta})) \right) + \theta_i \sum_{t=0}^{\infty} \delta^t \log e_{it}^*(\tilde{\theta}) \\
&\quad + (1 - \theta_i) \sum_{t=0}^{\infty} \delta^t \log(1 - \mathcal{E}_t^*(\tilde{\theta})) \\
&= \frac{\log \omega_0}{1 - \delta\gamma} + \sum_{t=0}^{\infty} \delta^t \sum_{j=0}^t \gamma^{t-j} \log(1 - \mathcal{E}_j^*(\tilde{\theta})) - \theta_i \sum_{t=0}^{\infty} \delta^t \log \left(\frac{1 - \mathcal{E}_t^*(\tilde{\theta})}{e_{it}^*(\tilde{\theta})} \right) \\
&\equiv r_i(\tilde{\theta}) - \theta_i q_i(\tilde{\theta})
\end{aligned} \tag{29}$$

According to this definition, for any type profile θ ,

$$r_i(\theta) \equiv \frac{\log \omega_0}{1 - \delta\gamma} + \sum_{t=0}^{\infty} \delta^t \sum_{j=0}^t \gamma^{t-j} \log(1 - \mathcal{E}_j^*(\theta)) \quad \text{and} \tag{30}$$

$$q_i(\theta) \equiv \sum_{t=0}^{\infty} \delta^t \log \left(\frac{1 - \mathcal{E}_t^*(\theta)}{e_{it}^*(\theta)} \right). \tag{31}$$

The payoff has a simple interpretation. $r_i(\theta)$ is the long run benefit of resource conservation given type profile θ reported to the IA. 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

$$R_i(\theta_i) \equiv \int_{\theta_{-i}} r_i(\theta) dF_{-i}(\theta_{-i}) \quad \text{and} \quad Q_i(\theta_i) \equiv \int_{\theta_{-i}} q_i(\theta) dF_{-i}(\theta_{-i}) \tag{32}$$

Hence, by definition,

$$\int U_i(\omega_0, \mathbf{c}^*(\theta), \theta) dF_{-i}(\theta_{-i}) \equiv R_i(\theta_i) - \theta_i Q_i(\theta_i) \tag{33}$$

The truth-telling constraint (27) may therefore be expressed as

$$R_i(\theta_i) - \theta_i Q_i(\theta_i) \geq R_i(\tilde{\theta}_i) - \theta_i Q_i(\tilde{\theta}_i) \quad \forall \theta_i \quad \forall \tilde{\theta}_i \tag{34}$$

Standard arguments show that any R_i and Q_i will satisfy (34) only if Q_i is weakly decreasing and

$$R_i(\theta_i) - \theta_i Q_i(\theta_i) = R_i(\bar{\theta}) - \bar{\theta} Q_i(\bar{\theta}) + \int_{\theta_i}^{\bar{\theta}} Q_i(\tilde{\theta}_i) d\tilde{\theta}_i \quad (35)$$

Equation (35) is the standard envelope condition. As is standard, the IC constraint for the highest type $\bar{\theta}$ does not bind.

Note that a standard integration by parts argument applied to (35) yields

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

Hence, any solution to the planner's problem solves

$$\max_{R_i(\bar{\theta}), Q_i(\cdot)} \sum_i \left[R_i(\bar{\theta}) - \bar{\theta} Q_i(\bar{\theta}) + \int_{\underline{\theta}}^{\bar{\theta}} F_i(\theta_i) Q_i(\theta_i) d\theta_i \right] \quad (36)$$

subject to Q_i weakly decreasing and (30), (31), (32), and (35).

Observe, first, that any solution to (36) must be stationary. To see this, observe from (30), (31), and (32) that all the components in (36) are discounted sums of log functions, and moreover, the initial state ω_0 appears only in $R_i(\theta)$ and as a constant. By dynamic consistency of the IA's optimal plan, the solution will be stationary.

Second, observe that in any solution \mathbf{e}^* to (36), $e_i^*(\theta) \leq \bar{\theta}(1 - \delta\gamma)/n$. To see why, observe that $\bar{\theta}(1 - \delta\gamma)/n$ is the full information optimum for the highest type $\bar{\theta}$ whose IC constraint does not bind. Consequently, if for some set of countries J and some neighborhood of type profiles N , $e_i^*(\theta) > \bar{\theta}(1 - \delta\gamma)/n$ for all $\theta \in N$ and all $i \in J$, then welfare according to (36) can be increased while preserving the monotonicity of Q_i by reducing e^* on this neighborhood of profiles so that $e_i^*(\theta) = \bar{\theta}(1 - \delta\gamma)/n$ for $\theta \in N$, $j \in J$. The first part, $\sum_i [R_i(\bar{\theta}) - \bar{\theta} Q_i(\bar{\theta})]$, of (36) is clearly increased by the change since $\bar{\theta}(1 - \delta\gamma)/n$ is the planner's full information optimal rate for types $\bar{\theta}$. As for the second part $\sum_i \left[\int_{\underline{\theta}}^{\bar{\theta}} F_i(\theta_i) Q_i(\theta_i) d\theta_i \right]$, since each Q_i is decreasing in e^* , the second part of (36) must also increase by the reduction of e^* on the neighborhood N .

Hence, let \mathbf{e}^* denote a stationary quota with $e_i^*(\theta_i) \leq \bar{\theta}(1 - \delta\gamma)/n$ that solves the planner's problem and suppose, by contradiction, that \mathbf{e}^* is not compressed. Then for some i , e_i^* is strictly monotone (increasing or decreasing) over some interval of types.

Recall the full information optimal quota, denoted here by \mathbf{e}^{**} (to distinguish it from the candidate solution \mathbf{e}^*), is given by $e_i^{**}(\theta_i) = \theta_i(1 - \delta\gamma)/n$. Because \mathbf{e}^{**}

prescribes an extraction rate for i that is independent of others' types, it represents the solution to (28) if the IC constraints in (27) are dropped.

Define the compressed quota \bar{e} by

$$\bar{e}_i(\theta_i) \equiv \bar{e}_i = e_i^{**}(\bar{\theta}) = \bar{\theta}(1 - \delta\gamma)/n \quad \forall \theta_i \quad \forall i.$$

This quota prescribes the optimal extraction rates as if i 's type was the highest usage type $\bar{\theta}$. Obviously, \bar{e} satisfies the incentive constraints. Hence, observe that $e_i^*(\theta_i) \leq \bar{e}_i$ for θ_i and all i .

Define

$$\begin{aligned} BR_i(\theta_i, e_{-i}^*, \theta_{-i}) &= \arg \max_{e_i} U_i(\omega_0, \mathbf{e}^*(\theta), \theta_i) \\ BR_i(\theta_i, e_{-i}^*, F_{-i}) &= \arg \max_{e_i} \int_{\theta_{-i}} U_i(\omega_0, \mathbf{e}^*(\theta), \theta_i) dF_{-i} \end{aligned}$$

These are country i 's ex post and interim best response functions, respectively. We claim that the parameter restriction

$$\frac{n\underline{\theta}}{1 + (n-1)\underline{\theta}(1 - \delta\gamma)} > \bar{\theta}. \quad (37)$$

implies

$$BR_i(\theta_i, e_{-i}^*, F_{-i}) > e_i^* \quad \forall \theta_i \quad (38)$$

To see why, observe first that

$$BR_i(\theta_i, e_{-i}^*, F_{-i}) \in co \{ BR_i(\theta_i, e_{-i}^*, \theta_{-i}) : \theta_{-i} \in [\underline{\theta}, \bar{\theta}]^n \}$$

where co denotes the convex hull of the set $\{\cdot\}$. From (9),

$$BR_i(\theta_i, e_{-i}^*, \theta_{-i}) = \theta_i(1 - \mathcal{E}_{-i}^*(\theta_{-i}))(1 - \delta\gamma)$$

Hence, a lower bound for the interim best response function, $BR_i(\theta_i, e_{-i}^*, F_{-i})$, can be derived by replacing $\mathcal{E}_{-i}^*(\theta_{-i})$ with $\bar{\mathcal{E}}_{-i} = \frac{(n-1)}{n}\bar{\theta}(1 - \delta\gamma) \geq \mathcal{E}_{-i}^*(\theta_{-i})$. Equation (37) is precisely the condition that ensures that

$$\theta_i(1 - \bar{\mathcal{E}}_{-i})(1 - \delta\gamma) \geq \bar{\theta}(1 - \delta\gamma)/n.$$

But (37) in fact holds under the dispersion assumption $n > \bar{\theta}/\underline{\theta}$ if $\delta\gamma$ is close to one. Hence, sufficient patience (as measured by $\delta\gamma$) implies (38).

Therefore, since \mathbf{e}^* is not compressed, it must be strictly increasing (strictly decreasing) in a neighborhood $N_i \subset [\underline{\theta}, \bar{\theta}]$ such that

$$e_i^*(\theta_i) < BR_i(\theta_i, e_{-i}^*) \quad \forall \theta_i \in N_i$$

But in this case, \mathbf{e}^* fails to be incentive compatible since types $\theta_i \in N_i$ can increase their payoffs by mimicking a strictly larger types $\tilde{\theta}_i > \theta_i$ (or strictly smaller types $\tilde{\theta}_i < \theta_i$). We conclude that any solution \mathbf{e}^* to the planner's problem (28) must be fully compressed.

We now show that among all fully compressed extraction rates, $(1 - \delta\gamma)/n \int \theta_i dF_i$ is the optimal extraction rate for all i .

To show this, notice that if \mathbf{e}^* is a fully compressed quota system, then $Q_i(\theta_i) = Q_i(\underline{\theta})$ for all types θ_i . Now evaluate (35) at $\theta_i = \underline{\theta}$ and set $Q_i(\theta_i) = Q_i(\underline{\theta})$ for all θ_i as prescribed by full compression. This yields:

$$R_i(\underline{\theta}) - Q_i(\underline{\theta})\bar{\theta} = R_i(\bar{\theta}) - \bar{\theta} Q_i(\bar{\theta})$$

Substituting this back into (36) yields

$$\sum_i \left[R_i(\underline{\theta}) - Q_i(\underline{\theta}) \left(\bar{\theta} - \int_{\underline{\theta}}^{\bar{\theta}} F_i(\tilde{\theta}_i) d\tilde{\theta} \right) \right] \quad (39)$$

By a simple integration by parts, $\bar{\theta} - \int_{\underline{\theta}}^{\bar{\theta}} F_i(\tilde{\theta}_i) d\tilde{\theta} = \int_{\underline{\theta}}^{\bar{\theta}} \tilde{\theta}_i dF_i(\tilde{\theta}_i)$, the average type for country i . We use this fact, together with the fact that $e_i(\theta)$ does not vary with type, to rewrite the objective in (36) as

$$\max_{r_i(\underline{\theta}), q_i(\underline{\theta})} \sum_i \left[r_i(\underline{\theta}) - q_i(\underline{\theta}) \left(\int_{\underline{\theta}}^{\bar{\theta}} \tilde{\theta}_i dF_i(\tilde{\theta}_i) \right) \right] \quad (40)$$

subject to (30) and (31). Since the functions $r_i(\underline{\theta})$ and $q_i(\underline{\theta})$ are simply choice variables of the planner, the solution to (40) yields the full information optimum in which each individual's type is $\int \theta_i dF_i$. The solution to this full information problem is the first best quota when i 's type is $\int \theta_i dF_i$. Specifically, the solution is given by

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

This concludes the proof of Part 1 of the Proposition. ■

Proof of Part 2 of Proposition 2. Given the compressed quota system \mathbf{c}^* from part 1, we construct a PBE (σ, μ) which implements \mathbf{c}^* . The proof largely mimics the steps of Proposition 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 ,

$$V_i^{\tau-1} \geq \max_{\theta_i} \left\{ \frac{1 - \delta\rho}{1 - \delta\rho - \delta^2(1 - \rho)} \max_{e_i} \left[\left(\frac{1}{1 - \delta\gamma} - \theta_i \right) \log(1 - \mathcal{E}_{-i}^{\tau-1} + e_i) + \theta_i \log e_i \right] \right. \\ \left. + \frac{\delta^2}{1 - \delta\rho - \delta^2(1 - \rho)} \left[\left(\frac{1}{1 - \delta\gamma} - \theta_i \right) \log(1 - \mathcal{E}^\tau) + \theta_i \log e_i^\tau \right] \right\} \quad (41)$$

Hence, the difference between (41) and the analogous constraint (26) 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 (41) by choosing \mathcal{E}^τ 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 $\sigma_i(h^0, \tilde{\theta}, \theta_i) = \omega_0 e_i^*$, 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 Proposition 3 . First, recall the definitions of $r_i(\tilde{\theta})$ and $q_i(\tilde{\theta})$ and of R_i and Q_i given by equations (30), (31), and (32).

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

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

Observe that

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

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

$$R_i(\theta_i) - \theta_i Q_i(\theta_i) + \int_{\theta_{-i}} s_i(\theta) dF_{-i} \geq R_i(\tilde{\theta}_i) - \theta_i Q_i(\tilde{\theta}_i) + \int_{\theta_{-i}} s_i(\tilde{\theta}_i, \theta_{-i}) dF_{-i} \quad \forall \theta_i \quad \forall \tilde{\theta}_i \quad (42)$$

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

$$\begin{aligned}
& R_i(\tilde{\theta}_i, \theta_{-i}) - \theta_i Q_i(\tilde{\theta}_i, \theta_{-i}) + \int_{\theta_{-i}} s_i(\tilde{\theta}_i, \theta_{-i}) dF_{-i} \\
&= \int_{\theta_{-i}} [r_i(\tilde{\theta}_i, \theta_{-i}) - \theta_i q_i(\tilde{\theta}_i, \theta_{-i}) + s_i(\tilde{\theta}_i, \theta_{-i})] dF_{-i} \\
&= \int_{\theta_{-i}} \left\{ r_i(\tilde{\theta}_i, \theta_{-i}) - \theta_i q_i(\tilde{\theta}_i, \theta_{-i}) + \sum_{j \neq i} [r_j(\tilde{\theta}_i, \theta_{-i}) - \theta_j q_j(\tilde{\theta}_i, \theta_{-i})] \right\} dF_{-i} + \text{2nd term of } s_i
\end{aligned}$$

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

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

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

Proof of Part 1 of Proposition 4 .

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

Let \mathbf{e}^* denote a solution to the super-relaxed planner's problem,

$$\begin{aligned}
& \max_{\mathbf{e}^*} \sum_i \int_{\theta_0} U_i(\omega_0, e^*(\theta_0), \theta_{i0}) dF_0(\theta_0) \quad \text{subject to} \\
& \int_{\theta_{-it}} U_i(\omega_t, e^*(\theta^t), \theta_{it}) dF_{-i}(\theta_{-it} | \theta_{t-1}) \geq \int_{\theta_{-it}} U_i(\omega_t, e^*(\theta^t \setminus \tilde{\theta}_{it}), \theta_{it}) dF_{-i}(\theta_{-it} | \theta_{t-1}) \quad \forall \theta^t \forall \tilde{\theta}_{it} \forall i \forall t
\end{aligned} \tag{43}$$

Written recursively, country i 's interim expected payoff at the date t disclosure

stage along the equilibrium path determined by type-history θ^{t-1} is given by

$$\int_{\theta_{-it}} U_i(\omega_t, e^*(\theta^t), \theta_{it}) dF_{-i}(\theta_{-it}|\theta_{t-1}) = \int_{\theta_{-it}} \left\{ \frac{\log(1 - \mathcal{E}_i^*(\theta^t))}{1 - \delta\gamma} - \theta_{it} \log\left(\frac{1 - \mathcal{E}_i^*(\theta^t)}{e_{it}^*(\theta^t)}\right) + \delta \int_{\theta_{t+1}} U_i(\omega_{t+1}, e^*(\theta^{t+1}), \theta_{it+1}) dF(\theta_{t+1}|\theta_t) \right\} dF_{-i}(\theta_{-it}|\theta_{t-1})$$

Note that the payoff does not depend on the current state ω_t . In particular, the effect of current extraction on the future state is entered through the first term $\frac{\log(1 - \mathcal{E}_i^*(\theta^t))}{1 - \delta\gamma}$. We can therefore drop the state from the notation and write i 's payoff as $\int_{\theta_{-it}} U_i(e^*(\theta^t), \theta_{it}) dF_{-i}(\theta_{-it}|\theta_{t-1})$.

If the planner were able to start anew and redesign the quota from date t onward given the information θ^{t-1} , then the recursive payoff would be used to define the planner's problem at that date. However, since the planner must commit to the quota at date $t = 0$, the recursive problem must be defined from date 0, given only his initial (prior) information.

Since, by definition, only one-shot deviations are considered in the super-relaxed problem (43), we can nevertheless apply Bellman's Principle of Optimality to the recursive structure defined from the initial date $t = 0$. We apply it as follows: fix any date t and consider the super-relaxed problem at that date given the planner's information (at date 0). We show that if the optimal solution from date $t + 1$ onward are fully compressed in the revealed information from date $t + 1$ onward, then the optimal solution at t will be as well.

Formally, we prove

Claim 1 *Let $\{\mathbf{e}_\tau^*\}_{\tau=t+1}^\infty$ denote the planner's super-relaxed solution from dates $t + 1$ onward. Suppose that $\{\mathbf{e}_\tau^*\}_{\tau=t+1}^\infty$ is invariant to all type profiles $\theta_{t+1}, \theta_{t+2}, \dots$. Then the planner's super-relaxed solution $\{\mathbf{e}_\tau^*\}_{\tau=t}^\infty$ from t onward will be invariant to all type profiles $\theta_t, \theta_{t+1}, \dots$.*

prf of Claim. Fix a date t and \mathbf{e}^* as in the hypothesis of the Claim. In the remainder of the analysis, we ignore the notational dependence of e^* on the history θ^{t-1} prior to date t (as it plays no role in the Claim). We can therefore rewrite the recursive

payoff of country i as

$$\int_{\theta_{-it}} \left\{ \frac{\log(1 - \mathcal{E}_t^*(\theta_t))}{1 - \delta\gamma} - \theta_{it} \log \left(\frac{1 - \mathcal{E}_t^*(\theta_t)}{e_{it}^*(\theta_t)} \right) + \delta \int_{\theta_{t+1}} U_i(e^*(\theta_t), \theta_{it+1}) dF(\theta_{t+1}|\theta_t) \right\} dF_{-i}(\theta_{-it})$$

$$\int_{\theta_{-it}} \left\{ \frac{\log(1 - \mathcal{E}_t^*(\theta_t))}{1 - \delta\gamma} - \theta_{it} \log \left(\frac{1 - \mathcal{E}_t^*(\theta_t)}{e_{it}^*(\theta_t)} \right) + \delta \int_{\theta_{it+1}} U_i(e^*(\theta_t), \theta_{it+1}) dF(\theta_{it+1}|\theta_{it}) \right\} dF_{-i}(\theta_{-it})$$

Using the no-drift assumption, observe that $\theta_{it} = E[\theta_{it+1}|\theta_{it}] = E[\theta_{it+2}|\theta_{it}] = \dots$. We apply this fact to the recursive payoff above. Since all future types $\theta_{i\tau}$ $\tau > t$ enter multiplicatively into the date τ flow payoff, we recurse forward writing out i 's payoff as

$$\int_{\theta_{-it}} \sum_{\tau=t}^{\infty} \delta^{\tau-t} \sum_{j=t}^{\tau} \gamma^{\tau-j} \log(1 - \mathcal{E}_j^*(\theta_t)) - \theta_{it} \sum_{\tau=t}^{\infty} \delta^{\tau-t} \log \left(\frac{1 - \mathcal{E}_{\tau}^*(\theta_t)}{e_{i\tau}^*(\theta_t)} \right) dF_{-i}(\theta_{-it})$$

Notice that is precisely the payoff $R_i(\theta_{it}) - \theta_{it}Q_i(\theta_{it})$ (minus the constant term in ω_t) defined by (31), (30), (32) in the perfect persistence case! By applying the result from that Proposition, we establish the present claim. \blacksquare

As for the remainder of the result, we apply Bellman's Principle, recursing forward from date 0. Hence,

$\{\mathbf{e}_t^*\}_{t=0}^{\infty}$ will be fully compressed if $\mathbf{e}_1^*, e_2^*, \dots$ are fully compressed in $\theta_1, \theta_2, \dots$

$\{\mathbf{e}_t^*\}_{t=1}^{\infty}$ will be fully compressed if $\mathbf{e}_2^*, e_3^*, \dots$ are fully compressed in $\theta_2, \theta_3, \dots$

\vdots

$\{\mathbf{e}_t^*\}_{t=\tau}^{\infty}$ will be fully compressed if $\mathbf{e}_{\tau+1}^*, e_{\tau+2}^*, \dots$ are fully compressed in $\theta_{\tau+1}, \theta_{\tau+2}, \dots$

\vdots

Thus, we obtain that the solution \mathbf{e}^* to the super-relaxed planner's problem in (43) is fully compressed. Observe, however, a fully compressed quota satisfies the full incentive constraints when multi-period deviations (in disclosure) are considered. The full Proposition follows if we can establish part 2. \blacksquare

References

- [1] Acemoglu, D., P. Aghion, L. Bursztyn, and D. Hemous (2012), “The Environment and Directed Technical Change,” *American Economic Review*, 102: 131-166.
- [2] Amador, M., I. Werning, and G.-M. Angeletos (2003): “Commitment Vs. Flexibility,” NBER Working Paper.
- [3] Aoyagi, M. (2003), “Bid Rotation and Collusion in Repeated Auctions,” *Journal of Economic Theory*, 112: 79-105.
- [4] Athey, S. and K. Bagwell (2008), “Collusion with Persistent Cost Shocks,” *Econometrica*, 76: 493-540.
- [5] Athey, S., K. Bagwell, and C. Sanchirico (2001), “Collusion and Price Rigidity,” *The Review of Economic Studies*, 17: 317-49.
- [6] 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.
- [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] Barrett, S. (2003), *Environment and Statecraft: The Strategy of Environmental Treaty-Making*, Oxford University Press, Oxford.
- [9] Battaglini, M. and Bard Harstad (2012), “Participation and Duration of Environmental Agreements,” NBER working paper, #18585.
- [10] Bodansky, D. (2004), “International Climate Efforts Beyond 2012: A Survey of Approaches,” *PEW Center on Global Climate Change*.
- [11] Burke, M., J Dykema, D. Lobell, E. Miguel, and S Satyanath (2011), “Incorporating Climate Uncertainty into Estimates of Climate Change Impacts, with Applications to U.S. and African Agriculture,” NBER Working paper #17092.
- [12] Calcott, P. and V. Petkov (2012), “International Environmental Agreements: Design of Optimal Transfers Under Heterogeneity,” *Environmental Modeling and Assessment*, 17: 209-220.
- [13] Cave, J. (1987), “Long-Term Competition in a Dynamic Game: The Cold Fish War,” *Rand Journal*, 18: 596-610.
- [14] d’Aspremont, C., and L.A. Gerard-Varet (1979), “Incentives and incomplete information,” *Journal of Public Economics*, 11: 25- 45.
- [15] Desmet, K. and E. Rossi-Hansberg (2013), “On the Spatial Economic Impact of Global Warming,” mimeo, Princeton University, February.

- [16] 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.
- [17] Dutta, P. and R. Radner (2006), "A Game-Theoretic Approach to Global Warming," *Advances in Mathematical Economics*, 8: 135-53.
- [18] Dutta, P. and R. Radner (2004), "Self-enforcing climate-change treaties," *Proceedings of the National Academy of Sciences*.
- [19] Finus, M. (2001), *Game Theory and International Environmental Cooperation*, Cheltenham, Edward Elgar.
- [20] Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski (2011), "Optimal Taxes on Fossil Fuel in General Equilibrium," NBER working paper.
- [21] Halac, M. and P. Yared (2012), "Fiscal Rules and Discretion under Persistent Shocks", NBER working paper.
- [22] Harrison, R. and R. Lagunoff (2013), "Business-as-Usual in the Global Commons," in progress.
- [23] Haurie, A. (2008), "A Dynamic Game Model of Strategic R&D Cooperation and GHG Emission Mitigation," mimeo.
- [24] Krusell, P. and A. Smith (2009), "Macroeconomics and Global Climate Change: Transition for a Many-Region Economy," working paper.
- [25] Levhari and L. Mirman (1980), "The Great Fish War: An Example Using a Dynamic Cournot-Nash Solution," *Bell Journal of Economics*, 14: 322-34.
- [26] N-V Long (2011), "Dynamic Games in the Economics of Natural Resources: A Survey," *Dynamic Games and Applications*, 1: 115-148.
- [27] McAfee, R.P. and J. McMillan, "Bidding Rings," *American Economic Review*, 82: 579-99.
- [28] Myerson, R. (1981), "Optimal Auction Design," *Mathematics of Operations Research*, 6: 58-73.
- [29] Nordhaus, W. (2006), "After Kyoto: Alternative Mechanisms to Control Global Warming," *American Economic Review*, 96: 31-34.
- [30] Nordhaus, W. (2007), "To Tax or Not to Tax: The Case for a Carbon Tax," *Review of Environmental Economics and Policy*.
- [31] Ostrom, E. (1990), *Governing the Commons: The Evolution of Institutions for Collective Action*, Cambridge University Press, New York.
- [32] Ostrom, E. (2002), "Common Pool Resource and Institutions: Towards a Revised Theory," *Handbook of Agricultural Economics, Volume 2*.
- [33] Pavan, A., I. Segal, and J. Toikka (2012), "Dynamic Mechanism Design," Stanford University mimeo.

- [34] Rouillon, S. (2010), "Optimal decentralized management of a natural resource," mimeo, Universite Bordeaux.
- [35] Roberts, K. (1984), "The Theoretical Limits of Redistribution," *Review of Economic Studies*, 51(2): 177-95.
- [36] Skyrpazc A. and H Hopenhayn (2004), "Tacit Collusion in Repeated Auctions," *Journal of Economic Theory*, 114: 153-69.
- [37] Stern, N. (2006), *The Stern Review Report on the Economics of Climate Change*.
- [38] United Nations Inter-Governmental Panel on Climate Change (2007), *Fourth Assessment Report: Climate Change*.