

Imperfect Monitoring May Be Perfectly Fine: Differential Game Theory and an Application to Climate Change

Stergios Athanassoglou *

March 11, 2009

Abstract

We study a class of finite-horizon differential games with control-state separable objective functionals and state equations. In this context, we introduce simple mechanisms that give a dynamic payoff to agents that only depends on the current value of the state variable. Thus, in many applications their implementation requires moderate monitoring capability. We are able to show that these mechanisms induce, in Markov-perfect equilibrium, any economically meaningful control path that is continuously differentiable. Moreover, if the target control-state path represents a Pareto-improvement over an inefficient unregulated outcome, a suitably chosen ex-ante cost-sharing scheme ensures ex-post individual rationality, in addition to budget balance. At the same time, if the differential game has pre-existing linear structure no such cost sharing agreement is necessary. We apply our analysis to a simple differential-game model of global climate change and exhibit a mechanism, which uniquely induces the globally optimal emissions path.

Keywords: differential games, imperfect monitoring, mechanism design, climate change, Volterra integral equations

JEL Classifications: C72, C73, Q53

*Post-Doctoral Fellow, The Earth Institute at Columbia University, New York, NY; e-mail: sa2164@columbia.edu. The author thanks Prajit Dutta, Ram Fishman, Glenn Sheriff, Rishi Talreja, and Anastasios Xepapadeas for helpful discussions and comments. The generous financial support of the Columbia Earth Institute Fellows Program is gratefully acknowledged.

1 Introduction

Motivation. The main goal of this paper is to shed light on the ability of simple policy tools to affect equilibrium behavior in a useful class of differential games. Importantly, the implementation of the proposed economic instruments does not require onerous monitoring. Indeed, monitoring requirements on agents' actions pose a significant challenge in the design and implementation of policy, under both noncooperative and cooperative assumptions on agents' behavior.

The same problems resonate in both cases, but let us consider the noncooperative framework first. In a world with unlimited monitoring capability, centralized policy makers may effectively attempt to influence equilibrium behavior through dynamic taxation (or subsidy) schemes on agents' direct actions. However, in most real-life settings this approach can be impractical, or even possibly misguided, for a number of reasons. As a purely technical matter, it requires strict and continuous tracking of agents' behavior which, in many cases, the policy maker may simply not be able to perform. This point is particularly relevant in instances where there are a great number of individual agents and monitoring devices are costly to install and maintain. A different, but equally important, problem arises from the incentives that such an approach introduces. If monitoring must be done by the agents themselves, and devices are not completely tamper-proof, then agents have strong incentives to misrepresent the data.

An alternative approach involves a cooperative game-theoretic framework. Assuming the absence of a centralized policy-making body, this approach is predicated on the notion that individual stakeholders come together to craft mutually binding agreements.¹ These agreements typically propose a cooperative course of action, and go on to stipulate punishments that deviating parties endure from the rest of the coalition for their misbehavior. Naturally, the discussion centers around designing these agreements so that agents prefer to follow the cooperative consensus to incurring the punishment of deviation. Much elegant analysis is done in determining credible "trigger" strategies that induce agents to improve upon the inefficient competitive outcome. Nonetheless, these strategies, once again, depend on agents actually observing their peers' deviations from the mutually agreed-upon control paths and are plagued by the same monitoring problems.

The above discussion illustrates that designing a cooperative or noncooperative incentive scheme that relies on monitoring agents' actions is problematic. However, while central policy makers and agents may not be able to observe individual decisions by their constituents or peers, it is likely that they are able to monitor the state variable. This is because, in many prac-

¹See Barrett [2] for a comprehensive account of this literature in the context of environmental treaty-making.

tical cases, the state is of much lower dimension than the control; as a result, fewer monitoring devices are necessary, driving down the cost of policy implementation. Furthermore, these devices would be harder to tamper with as agents would, typically, not have easy access to them.² Appreciating the soundness of these arguments, our work is concerned with designing effective policy instruments that only take into account the current value of the relevant state variables, under a noncooperative framework.

Related Work. Differential games are used in a diverse set of economic contexts. Indeed, the theory of optimal control proves to be a powerful tool in deriving insights into a great number of applications spanning industrial organization, R and D and capital accumulation [17, 18, 8], public economics [13, 21], and environmental and resource economics [22, 15, 14, 16].

Our work is influenced by an early contribution due to Xepapadeas [22]. Focusing on a model of dynamic nonpoint source pollution –which is plagued by the kinds of monitoring problems discussed in the previous section– he makes a strong case for so-called “ambient” taxes on the pollution stock. These would reward or penalize agents solely on the basis of pollution levels’ deviation from a socially optimal path. Xepapadeas goes on to exhibit an ambient taxation scheme that, in steady state, induces socially desirable emissions; moreover, he makes the analysis robust by introducing uncertainty into the model. However, one important point he does not raise is *how* this steady state is reached. In particular, issues of potential inter-temporal welfare loss (in relation to the social optimum) en route to the equilibrium are not explored. Our work addresses this exact issue in finite-horizon problems and shows that, if a mechanism is chosen appropriately, no welfare loss will occur.

In recent years, the issues of time-consistency and credibility of incentive mechanisms have received attention in the dynamic game literature [15, 14, 9, 16, 9, 11]. A common thread running through many of these contributions is their focus on games which admit a special “linear-state” structure (see Chapter 7.2 in Dockner et al. [7]). This consistency is not entirely coincidental; linear state games have the appealing property that their open-loop Nash equilibria are Markov-perfect (Dockner et al. [6], Fershtman [12]) and therefore admit meaningful analysis.³ We comment on this literature in a little more depth since it (a) deals with linear state games, which play a key

²Consider the example of emissions control. Agents can easily tamper with, say, a factory’s electricity meter; doing the same with a pollution-monitoring device that is suspended in the atmosphere is not quite as easy. Xepapadeas [22] elaborates on this point.

³Linear state games are closely related to exponential games, which have been used in the R and D literature by Reinganum [17, 18].

role in our subsequent analysis and; (b) features a “mechanism design” framework for inducing agents to adopt mutually beneficial behavior that is, in spirit if not in methodology, congruent to the one presented in our work.

An incentive mechanism, agreed upon jointly by agents in the beginning of a game, is said to satisfy time-consistency if it induces desirable equilibrium behavior that is subgame-perfect. In this context, Jorgensen et al. [14] identify a necessary condition that linear-state games must satisfy if their cooperative outcomes are to be time-consistent; when this condition is not satisfied they describe a set of side-payments, which were also discussed in earlier work by the same authors [15], that may be used to ensure the desired result. In follow-up work, Martin-Herran and Zaccour [16] study, in a two-player linear-state setting, incentive equilibrium strategies. These are, essentially, trigger strategies that an agent will revert to if his adversary deviates from the mutually agreed-upon solution. In this context, they provide complex necessary conditions for a strategy pair to be an incentive equilibrium at the cooperative solution. They further focus on the credibility of incentive strategies, a property that requires each player to stick to the agreed-upon incentive strategy and not revert to the cooperative solution, even when one of his peers chooses to break the agreement.

In a series of recent papers, Dutta and Radner [9, 10, 11] study a linear-state dynamic game of global climate change. Their main goal is to characterize the kinds of equilibrium behavior that may be sustained under “trigger” threats that countries cooperatively agree to enforce if one of their peers violates the terms of a joint agreement (i.e., a climate change treaty). To this end, they describe the first-best outcome (a global Pareto optimum) and design equilibrium strategies that improve upon the inefficient status quo equilibrium (which they refer to as the “business-as-usual” equilibrium) that is defined by a tragedy-of-the-commons nature. An interesting property of their model is that it gives rise to Markov-perfect equilibrium emissions that are constant over time. Moreover, Dutta and Radner are able to incorporate issues of population growth and technological change into their analysis [10]. Our work in Section 5 represents a methodological shift from their framework, in a way to be made clear shortly.

Contribution. We focus on a class of finite-horizon differential games that is control-state separable, but allows for nonlinearities in the objective functionals and state equations. Having set up the scope of our inquiry, we introduce a class of “linear-state” mechanisms, which cancel out agents’ original cost functions and replace them with a dynamic payoff that is (unsurprisingly) linear in the *state*. This kind of policy intervention has the effect of giving the game a linear-state structure that allows for Markov-perfect, and therefore insightful, equilibrium analysis. In

contrast to [15, 14, 16, 9, 11] our framework is noncooperative. We further emphasize that the proposed mechanisms do not take the *control* into account when determining agents' payoffs; this point is important for two reasons. First, compared to standard dynamic taxation or subsidy schemes on agents' direct actions, these mechanisms are *de facto* more constrained in influencing equilibrium behavior. Thus, getting traction out of them –as a technical matter– is theoretically challenging. Second, along the lines outlined earlier, focusing on mechanisms that keep track of the state variable is desirable in instances where there are insufficient resources to reliably monitor or truthfully elicit agents' individual actions.

Our paper's main contribution is a general existence result regarding the power of linear-state mechanisms in the class of differential games we introduce. In particular, we show that for any continuously differentiable and rational⁴ control path, there exists a linear-state mechanism that induces exactly that path in Markov-perfect equilibrium. Moreover, this mechanism can be implemented with an ex-ante cost-sharing scheme that ensures budget balance and, when applied to target paths that represent a Pareto improvement over an unregulated status quo outcome, ex-post individual rationality. In fact, for games with pre-existing linear structure cost-sharing agreements are, in this regard, not necessary. To the best of our knowledge, our work departs from previous contributions in that it simultaneously: (a) addresses cost and state equation nonlinearities; (b) deals with a finite time horizon and multi-dimensional state space; (c) introduces first-best incentive schemes while discussing issues of budget balance and individual rationality; (d) adopts a noncooperative framework that does not presuppose agents' joint coordination; (e) proposes policy tools that do not require monitoring of agents' controls and (f) results, in many cases,⁵ in simple closed-form policy prescriptions. Furthermore, our proof technique relies on the theory of Volterra integral equations and, we believe, may be of independent interest in determining policy tools for tractable classes of differential games.

We demonstrate the applicability of our analysis with a simple differential game of climate change analyzed by Dutta and Radner [9, 10, 11]. Assuming that the world agrees to play a noncooperative game administered by some multinational body⁶ (i.e., the U.N.), we explicitly

⁴Essentially, a path is rational if it is feasible and, in some cases, not obviously suboptimal in the original unregulated environment.

⁵These typically correspond to problem instances with a low-dimensional state space. For an example, see Section 5 of this paper, as well as the adjacent-aquifer groundwater game studied in Athanassoglou et al. [1].

⁶We concede that this represents a substantial departure from Dutta and Radner's framework as it (optimistically) presupposes the existence of an effective world body that can enforce the rules of the stipulated noncooperative game. And while in many economic contexts the assumption of a game administrator is appropriate (consider, for instance, the mechanism-design literature on auctions), reservations about applying it to climate-change policy are

compute the linear-state mechanism, which, in Markov-perfect equilibrium, uniquely induces the globally optimal emissions path. Important departures of our model from Dutta and Radner involve our treatment of nonlinear cost functions and adoption of a finite time horizon. With regard to the latter, we view the assumption of a finite horizon as a useful one, especially because climate change agreements involve a relatively long, but still finite, timeframe. In fact, it can be argued that effective policy can only result from short to medium-term agreements.⁷ Moreover, since the target emissions path represents a Pareto-improvement over the status quo, it is possible to institute an ex-ante cost-sharing arrangement (possibly implemented by an agency like the World Bank) that ensures ex-post individual rationality, and therefore guarantees voluntary participation in the mechanism. Finally, if agents' state-dependent costs are linear, as modeled in Dutta and Radner's work, no ex-ante cost-sharing scheme would be necessary for our mechanism to induce the global optimum in Markov-perfect equilibrium.

Structure of the Paper. The rest of the paper is organized as follows. Section 2 introduces the model and the class of problems we will be focusing on. Section 3 describes linear-state mechanisms and derives the necessary and sufficient equilibrium conditions they give rise to. Section 4 contains a proof of the general existence result and discusses issues of budget balance and individual rationality. Section 5 applies our analysis to a differential game of global climate change. We provide concluding remarks in Section 6.

2 Model Description

Consider a n -player differential game where the variable $q_i(t) \in \mathfrak{R}$ denotes agent i 's control and $x(t) \in \mathfrak{R}$ denotes the state at time t . The n -dimensional control vector is denoted by $q(t)$. A control variable q_i is assumed to be non-negative, i.e., $q_i \in \mathfrak{R}_+$, but our analysis easily extends to the more general case where

$$q_i(t) \in [Q_i^1, Q_i^2] \text{ where } -\infty \leq Q_i^1 \leq Q_i^2 \leq \infty, \text{ for } i \in \{1, 2, \dots, n\}.$$

We assume that the payoff function of agent i is given by the expression

$$h_i(q(t), t) + c_i(x(t), t),$$

justified.

⁷What is the point of countries earnestly vowing to achieve ambitious goal x by the year 2100? Is such a long time horizon meaningful in designing policy right now? While theoretically appealing, the elegant steady-state analysis pursued by many authors may, in some cases, conflict with basic intuition. Physical conditions, technological progress and, perhaps most importantly, political winds are unpredictable.

where $h_i(\cdot)$ is strictly concave in q_i and twice continuously differentiable, and $c_i(\cdot)$ is a cost function, whose structure will be described in Assumption 1. Moreover, the state variable evolves according to the differential equation

$$\dot{x}(t) = f(q(t), t) + g(t)x(t),$$

where $f(\cdot)$ is a twice continuously differentiable function whose structure will be described in Assumptions 1 and 2. The function $g(\cdot)$ is continuously differentiable. Finally, the terminal time state-dependent cost of agent i is given by a function

$$K_i(x(T)),$$

where $K_i(\cdot)$ will, again, be discussed in Assumption 1.

We make the following assumptions on our model's structure.

Assumption 1 (*Structure*) *The functions $c_i(\cdot)$, $f(\cdot)$, and $K_i(\cdot)$ satisfy one of the following conditions for all $i \in \{1, 2, \dots, n\}$.*

- (a) *The function $f(q; t)$ is linear in q_i and separable in q_i and q_{-i} , and $c_i(x; t), K_i(x)$ are arbitrary, for each $t \in [0, T]$.*
- (b) *The function $f(q; t)$ is strictly convex-increasing in q_i , and $c_i(x; t), K_i(x)$ are decreasing in x , for each $t \in [0, T]$. Moreover, we assume that*

$$\frac{\partial}{\partial q_i} h_i(q; t) \geq 0, \text{ for all } q \text{ such that } q_i = 0, \text{ and each } t \in [0, T].$$

Part (a) of Assumption 1 does not require much discussion. Assuming $f(q; t)$ is linear in q_i and (q_i, q_{-i}) -separable, it does not impose any further requirements on the model. On the other hand, Part (b) may call for some explanation and economic context. Its prescriptions are typically satisfied in instances in which the state variable x corresponds to a costly by-product of economic activity (denoted by the control variable q), such as pollution. Moreover, an increase in this economic activity is assumed to produce increased damages. In this context, the requirement on h_i 's partial derivative is made to rule out trivial cases in which an agent would always set his control to zero. The strict convexity of $f(\cdot)$ in q_i is imposed for technical reasons that will become apparent in Section 3.2 and the subsequent proof of Theorem 1.

Assumption 2 (*Non-Triviality*) *Agents always have a non-trivial marginal effect on the state variable, i.e.,*

$$f_{q_i}(q, t) \neq 0, \text{ for all } q \text{ and } t \in [0, T] \text{ and } i \in \{1, 2, \dots, n\}$$

This assumption implies that agent i 's control always has *some* marginal effect on the state. Imposing this requirement is crucial because it enables us to design mechanisms, which focus on and exploit this exact dependence to affect equilibrium behavior. Following Theorem 1, we discuss how this assumption can be relaxed so that it applies only for control paths in which agent i 's control is interior.

While it is not our intention to oversell the reach of our modeling assumptions, we close this section by noting that they are satisfied in a number of prior contributions. For example, simple linear instances of our model may be found in Fershtman and Nitzan [13], Xepapadeas [22], Dockner et al. [8], Wirl [21], Jorgensen et al. [14], Martin-Herran and Zaccour [16], and Dutta and Radner [9, 10, 11]. One could undoubtedly find more references to this effect, so the above discussion is not meant to be exhaustive, merely indicative. From our point of view, we believe that our framework is especially appropriate for many problems, which naturally arise in environmental and resource economics.

3 Mechanism Design

3.1 Linear-State Mechanisms

In this section we discuss policy instruments for the class of differential games we have introduced. We focus our attention on the intuitive family of additive mechanisms, whose effect on agents' strategic behavior can be described in simple terms.

Definition 1 *Suppose $x(t) \in X$ for all $t \in [0, T]$, where $X \subseteq \mathfrak{R}$. An additive mechanism ϕ is a set of functions $\phi_i : X \times [0, T] \mapsto \mathfrak{R}$, for $i \in \{1, 2, \dots, n\}$.*

Note that, by definition, we do not allow mechanisms to depend on the agents' controls. Each additive mechanism ϕ induces a differential game between the agents. Focusing on agent i , and denoting other agents controls by q_{-i}^* , ϕ gives rise to the following differential game:

$$\begin{aligned} \max_{q_i(\cdot)} \quad & \int_0^T e^{-\delta t} [h_i(q_i(t), q_{-i}^*(t), t) + c_i(x(t)) + \phi_i(x(t), t)] dt + e^{-\delta T} (\phi_i(x(T), T) + K_i(x(T))) \\ \text{subject to:} \quad & \dot{x}(t) = f(q_i(t), q_{-i}^*(t), t) + g(t)x(t) \\ & q_i \geq 0, \quad x(0) = \hat{x}. \end{aligned} \tag{1}$$

In what follows we give a formal definition of a Markovian-Nash equilibrium for (1).

Definition 2 Suppose $x(t) \in X$ for all $t \in [0, T]$, where $X \subseteq \mathfrak{R}$. A set $(q_1^*, q_2^*, \dots, q_n^*)$ of functions where $q_i^* : X \times [0, T] \mapsto \mathfrak{R}$, is called a Markovian-Nash equilibrium if, for each $i \in \{1, 2, \dots, n\}$ an optimal control path $q_i(\cdot)$ of the maximization problem given by (1) exists and is given by the Markovian strategy $q_i(t) = q_i^*(x(t), t)$.

As we can see, a Markovian-Nash equilibrium is defined by Markovian strategies that take into account the current state and calendar time. Its definition equips us to introduce a stronger equilibrium concept, which we will be focusing on in this paper.

Definition 3 A Markov-perfect equilibrium is a subgame-perfect Markovian-Nash equilibrium.

If we restricted ourselves to open-loop strategies, in which agents pre-commit to an entire control path in the beginning of the time horizon, we obtain a different sort of equilibrium.

Definition 4 A set $(q_1^*, q_2^*, \dots, q_n^*)$ of functions where $q_i^* : [0, T] \mapsto \mathfrak{R}$, is called an open-loop Nash equilibrium if, for each $i \in \{1, 2, \dots, n\}$ an optimal control path $q_i(\cdot)$ of the maximization problem given by (1) exists and is given by the open-loop strategy $q_i(t) = q_i^*(t)$.

By definition, open-loop equilibria are degenerate Markovian-Nash equilibria; however, they typically do not satisfy subgame-perfectness. In fact, Markov-perfect equilibria are far more compelling than their open-loop counterparts since they allow agents to adapt their strategies to changes in the state variable –in a way that maximizes their payoff-to-go, even when they may find themselves off the equilibrium path. Indeed, restricting one’s attention to (non-subgame perfect) open-loop Nash equilibria can only be justified in instances where the state vector is not observable, rendering any such adaptive capabilities moot. At the same time, the attractive features of Markov-perfect equilibria come at a steep price as they are typically very hard to compute. Fortunately, this turns out not be to a problem in our case, as the class of differential games we study will reduce to one that admits tractable Markov-perfect analysis.

We focus our attention on an intuitive class of additive mechanisms, in which ϕ_i is the sum of a linear function of $x(t)$ and the negative of agent i ’s original $c_i(\cdot)$ function. We refer to this class of mechanisms as *linear-state* mechanisms. In particular, we choose ϕ_i so that

$$\phi_i(x(t), t) = \beta_i(t)x(t) - c_i(x(t)).$$

Complementing their mathematical tractability, linear-state mechanisms represent a policy scheme that is simple enough for real-life economic agents to understand. We further impose a terminal payoff along similar lines. In particular, the terminal time payoff is written as

$$\phi_i(x(T), T) = \gamma_i x(T) - K_i(x(T)).$$

Abusing notation slightly, given a problem instance $c(\cdot), K(\cdot)$, we refer to a linear-state mechanism ϕ such that

$$\begin{aligned}\phi_i(x(t), t) &= \beta_i(t)x(t) - c_i(x(t)), \quad t \in [0, T] \\ \phi_i(x(T), T) &= \gamma_i x(T) - K_i(x(T))\end{aligned}\tag{2}$$

for $i \in \{1, 2, \dots, n\}$ as a linear-state mechanism (β, γ) .

At this point, the introduction of the agents' original cost functions in the mechanism might give the reader pause, since it is reasonable to wonder how such a mechanism may be implemented to satisfy budget balance and individual rationality. We attempt to mitigate some of these concerns in Section 4.2. In any case, having described their structure, we are now ready to discuss linear-state mechanisms' equilibrium properties.

3.2 Linear-State Mechanism Equilibrium Derivation

Denoting other agents' strategies by q_{-i}^* , a linear-state mechanism (β, γ) applied to (1), induces the following, simpler, differential game:

$$\begin{aligned}\max_{q_i(\cdot)} \quad & \int_0^T e^{-\delta t} [h_i(q_i(t), q_{-i}^*(t), t) + \beta_i(t)x(t)] dt + e^{-\delta T} \gamma_i x(T) \\ \text{subject to:} \quad & \dot{x}(t) = f(q_i(t), q_{-i}^*(t), t) + g(t)x(t) \\ & q_i \geq 0, \quad x(0) = \hat{x}.\end{aligned}\tag{3}$$

The game given by (3) is a linear state game (see chapter 7.2 in Dockner et al. [7]). It is a well-known result that all open-loop equilibria of such games are Markov-perfect (see Dockner et al. [6], Fershtman [12], and pages 187-189 in Dockner et al. [7]), so we may note the following proposition without proof.

Proposition 1 *All open-loop Nash equilibria of differential game (3) are Markov-perfect.*

Proposition 1 is extremely important because it allows us to focus on deriving open-loop equilibria for this game, trusting that they correspond to subgame-perfect, and therefore meaningful, equilibrium behavior.

Thus, we proceed to calculate a Markov-perfect open-loop equilibrium of the differential game given by (3) induced by a linear-state mechanism (β, γ) . The current-value Hamiltonian of agent i is:

$$H_i(q_i, q_{-i}^*, x, \lambda_i, t) = h_i(q_i, q_{-i}^*, t) + \beta_i(t)x + \lambda_i \left(f(q_i, q_{-i}^*, t) + g(t)x \right).$$

We turn to the co-state variable and obtain the following differential equation

$$\dot{\lambda}_i(t) = (\delta - g(t))\lambda_i(t) - \beta_i(t). \quad (4)$$

We further impose the transversality condition

$$\lambda_i(T) = \gamma_i. \quad (5)$$

Solving (4) with terminal condition (5) and explicitly noting the co-state variable's dependence on (β_i, γ_i) yields

$$\lambda_i^{\beta_i, \gamma_i}(t) = \gamma_i e^{-\int_t^T (\delta - g(s)) ds} + \int_t^T \beta_i(s) e^{-\int_t^s (\delta - g(\tau)) d\tau} ds. \quad (6)$$

At this point, we need to consider the concavity of the Hamiltonian H_i with respect to (q_i, x) . To this end, recall that the function h_i is strictly concave in q_i and consider Assumption 1. If part (a) of Assumption 1 is satisfied, then H_i is clearly strictly concave in (q_i, x) . If part (b) is satisfied, then we restrict $\beta_i(\cdot)$ and γ_i so that $\lambda_i^{\beta_i, \gamma_i}(t)$ is negative for all $t \in [0, T]$. This, along with the stipulated strict convexity of $f(\cdot)$ in q_i , would again ensure that H_i is strictly concave in (q_i, x) . Having argued this point, the Hamiltonian-maximizing condition is given by

$$\begin{aligned} & \frac{\partial}{\partial q_i} h_i(q_i(t), q_{-i}^*(t), t) + \lambda_i^{\beta_i, \gamma_i}(t) f_{q_i}(q_i(t), q_{-i}^*(t), t) \leq 0 \\ q_i(t) \left(\frac{\partial}{\partial q_i} h_i(q_i(t), q_{-i}^*(t), t) + \lambda_i^{\beta_i, \gamma_i}(t) f_{q_i}(q_i(t), q_{-i}^*(t), t) \right) &= 0, \quad q_i(t) \geq 0. \end{aligned} \quad (7)$$

As the Hamiltonian of agent i is strictly concave in (q_i, x) and his terminal payoff is linear (and therefore concave) in x , conditions (6) and (7) for all $i \in \{1, 2, \dots, n\}$ characterize a set of Markov-perfect open-loop equilibria (see Sethi and Thompson [20]). To provide an explicit reference for the assertion of Proposition 1 we note the following result due to Fershtman [12].

Proposition 2 (*Lemma 3.4 in Fershtman [12]*) *Consider a differential game G . If the Pontryagin-type necessary conditions for an open-loop Nash equilibrium do not depend on the state variables, then an open-loop equilibrium, if it exists, is a degenerate Markov-perfect equilibrium.*

Clearly, the Pontryagin-type necessary (and sufficient) conditions for an open-loop equilibrium given by Expressions (6) and (7) do not depend at all on the state variable. Thus, Proposition 1 is valid.

4 An Existence Result

4.1 Proof of the Main Theorem

We begin this section by defining control paths that satisfy a reasonable criterion of rationality.

Definition 5 *A control path \hat{q} is rational if the following statements hold for all $i \in \{1, 2, \dots, n\}$.*

(i) *If agent i satisfies part (a) of Assumption 1, then \hat{q}_i is feasible.*

(ii) *If agent i satisfies part (b) of Assumption 1, then \hat{q}_i is feasible and satisfies*

$$\frac{\partial h_i(\hat{q}_i, \hat{q}_{-i}, t)}{\partial q_i} \geq 0, \quad \text{all } t \in [0, T].$$

The first statement of Definition 5 does not require much discussion; under separability and linearity of the state equation in q_i , any feasible control path for agent i is considered rational. The second statement is meant to restrict attention to paths that are economically meaningful in the following sense. Recall that under part (b) of Assumption 1, $f(\cdot)$ is strictly convex-increasing in q_i , and $c_i(\cdot), K_i(\cdot)$ are decreasing in x . In this kind of economic environment, given his peers' controls, an agent i would have no rational reason to set his control to a value q_i that results in a strictly negative partial derivative of h_i with respect to q_i , as doing so is clearly suboptimal. This is because it is easy to see that x is strictly increasing in the control q_i ; thus, setting q_i beyond the point where h_i 's partial (with respect to q_i) becomes zero results in losses both in terms of the function h_i as well as the functions c_i and K_i . In other words, this kind of control path will never be adopted by a profit-maximizing agent i . Thus, agent i will always set his control to \hat{q}_i small enough so that

$$\frac{\partial h_i(\hat{q}_i, \hat{q}_{-i}, t)}{\partial q_i} \geq 0 \quad \text{for all } t \in [0, T].$$

By the last statement in part (b) of Assumption 1, this is possible, since we have imposed that

$$\frac{\partial}{\partial q_i} h_i(q, t) \geq 0, \quad \text{for all } q \text{ such that } q_i = 0, \text{ and } t \in [0, T].$$

In particular, a rational path for agent i will always exist.

We proceed to show that linear-state mechanisms have the ability to induce, in Markov-perfect open-loop equilibrium, any continuously differentiable rational control path over $[0, T]$. Given a linear-state mechanism (β, γ) , let $\mathbf{Q}^{\beta, \gamma}$ denote the set of Markov-perfect open-loop equilibrium control paths over $[0, T]$ satisfying Expressions (6) and (7) for all $i \in \{1, 2, \dots, n\}$.

Theorem 1 Suppose \hat{q} is a continuously differentiable and rational control path in $[0, T]$. There exists a linear-state mechanism (β, γ) such that

$$\hat{q} \in \mathbf{Q}^{\beta, \gamma}.$$

Proof. Given the necessary and sufficient Markov-perfect open-loop equilibrium conditions (6) and (7) it is sufficient to consider mechanisms (β, γ) that (attempt to) satisfy

$$\begin{aligned} \frac{\partial}{\partial q_i} h_i(\hat{q}(t), t) &= -f_{q_i}(\hat{q}(t), t) \left(\gamma_i e^{-\int_t^T (\delta - g(s)) ds} + \int_t^T \beta_i(s) e^{-\int_t^s (\delta - g(\tau)) d\tau} ds \right) \\ &\text{for all } i \in \{1, 2, \dots, n\} \text{ and } t \in [0, T]. \end{aligned} \quad (8)$$

This restriction is well-defined, since by definition, the target path \hat{q} that we wish to induce is rational.⁸

So we proceed to determine an appropriate mechanism. Focusing on an agent i , by Assumption 2, we have

$$f_{q_i}(\hat{q}(t), t) \neq 0, \text{ for all } t \in [0, T].$$

Thus, we can always divide by this quantity. To satisfy Equation (8) for $t = T$ we set

$$\hat{\gamma}_i = -\frac{\frac{\partial}{\partial q_i} h_i(\hat{q}(T), T)}{f_{q_i}(\hat{q}(T), T)}. \quad (9)$$

To satisfy Equation (8) for $t \in [0, T]$ it is necessary and sufficient to impose that $\beta_i(\cdot)$ solve the following integral equation

$$\frac{\frac{\partial}{\partial q_i} h_i(\hat{q}(t), t)}{f_{q_i}(\hat{q}(t), t)} + \hat{\gamma}_i e^{-\int_t^T (\delta - g(s)) ds} = -\int_t^T \beta_i(s) e^{-\int_t^s (\delta - g(\tau)) d\tau} ds. \quad (10)$$

Performing the change of variable $z = T - t$, and denoting the left-hand-side of (10) evaluated at $t = T - z$ by $F(z)$, the previous equation obtains

$$F(z) = \int_0^z \beta_i(T - s) e^{-\int_s^z (\delta - g(T - \tau)) d\tau} ds. \quad (11)$$

Equation (11) is a linear Volterra equation of the first kind, whose kernel

$$\Theta(z, s) = e^{-\int_s^z (\delta - g(T - \tau)) d\tau}$$

⁸First, the rationality of \hat{q} imposes that the induced equilibrium path be feasible. Second, it ensures that Expression (6) is negative for agents that satisfy part (b) of Assumption 1 (recall that, in this case, $f(\cdot)$ is strictly convex increasing in q_i), which is required for their equilibrium conditions (7) to be sufficient.

is such that

$$\Theta(z, z) = 1 \text{ for all } z \in [0, T].$$

As all the functions which appear in Equation (11) are continuously differentiable, we can conclude that (a) the left-hand-side of integral equation (11) is continuously differentiable in $[0, T]$ and; (b) its kernel is continuously differentiable in $[0, T] \times [0, T]$. All of the above observations imply that Equation (11) can be reduced to the following equivalent Volterra equation of the second kind:

$$\frac{F'(z)}{\Theta(z, z)} = \beta_i(T-z) + \int_0^z \frac{\Theta_z(z, s)}{\Theta(z, z)} \beta_i(T-s) ds \Rightarrow F'(z) = \beta_i(T-z) + \int_0^z \Theta_z(z, s) \beta_i(T-s) ds. \quad (12)$$

By Theorem 2.1.1 in Burton [3], Equation (12) has a unique solution $\hat{\beta}_i(\cdot)$, which can be computed via Picard's method of successive approximation.

Repeating this argument for agents $j \neq i$ and collecting all the functions $\hat{\beta}_k(\cdot)$ and constants $\hat{\gamma}_k$ for $k \in \{1, 2, \dots, n\}$ yields a mechanism $(\hat{\beta}, \hat{\gamma})$ satisfying (8). This completes the proof. \blacksquare

Remark. In light of Theorem 1, one may wish to impose additional structure on the functions h_i and f to ensure that the sets $\mathbf{Q}^{\beta, \gamma}$ are singletons. Doing so would help in dealing with a game's (potential) equilibrium non-uniqueness, which may result in undesirable policy outcomes. This sort of structure is present in the climate change game we will examine in Section 5.

4.2 Cost Sharing, Budget Balance, and Ex-Post Individual Rationality

In an economic environment where participation in any sort of policy framework is voluntary, it is important to consider how the costs of policy implementation are shared and whether agents find it in their best interest to participate in the proposed scheme. To this end, consider the state trajectory \hat{x} brought about by the target control path \hat{q} , and the subsequent modification of the mechanism charges $\hat{\beta}_i(t)x(t)$ and $\hat{\gamma}_i x(T)$ that appear in Expression (2) to

$$\hat{\beta}_i(t)(x(t) - \hat{x}(t)) \text{ and } \hat{\gamma}_i(x(T) - \hat{x}(T)).$$

As this modification merely adds a constant to agents' objective functions, Theorem 1 applies and we once again obtain \hat{q} (and therefore also \hat{x}) in Markov-perfect open-loop equilibrium. Moreover, agent i 's induced equilibrium payoff is given by

$$\int_0^T e^{-\delta t} \left(h_i(q^{\hat{\beta}, \hat{\gamma}}(t)) + \hat{\beta}_i(t) \underbrace{[x^{\hat{\beta}, \hat{\gamma}}(t) - \hat{x}(t)]}_0 \right) dt + e^{-\delta T} \hat{\gamma}_i \underbrace{(x^{\hat{\beta}, \hat{\gamma}}(T) - \hat{x}(T))}_0 = \int_0^T e^{-\delta t} h_i(\hat{q}(t)).$$

In particular, the monetary effect of the mechanism charges $\beta_i(\cdot)$ and γ_i vanishes. On the other hand, recalling that linear-state mechanisms, by definition (see Expression (2)), cancel out agents' original cost functions, the total cost of policy implementation is given by the following expression

$$\sum_{i=1}^n \left(\int_0^T e^{-\delta t} c_i(\hat{x}(t)) dt + e^{-\delta T} K_i(\hat{x}(T)) \right).$$

In order to ensure budget balance, this cost may be allocated in any number of ways among the n agents. At the same time, note that any potential agent cost-shares are, from the agents' point of view, constants only depending on the exogenously given target state path \hat{x} . In the case in which the target outcome path (\hat{q}, \hat{x}) represents an ex-post (meaning ex-post terminal time T) Pareto-improvement over an unregulated competitive outcome, the cost shares may be set ex-ante (i.e., at time $t = 0$) to a value of

$$\int_0^T e^{-\delta t} c_i(\hat{x}(t)) dt + e^{-\delta T} K_i(\hat{x}(T)) \quad \text{for } i \in \{1, 2, \dots, n\}.$$

Under such a framework, it is assumed that the (modification of the) mechanism $(\hat{\beta}, \hat{\gamma})$ is implemented *if and only if all agents pay their cost shares at the beginning of the time horizon*. With this arrangement in place, all agents, who wish to improve upon their total discounted payoff over $[0, T]$, will, in Nash equilibrium, voluntarily bear these ex-ante costs and participate in the mechanism. In particular, *any* kind of ex-ante cost sharing arrangement that ensures an ex-post Pareto improvement over the unregulated equilibrium would satisfy ex-post individual rationality and voluntary participation.

In a sense, these ex-ante cost shares represent the agents' credible commitment to working together in order to improve their individual, as well as common, lot. They may also be thought of as a bank deposit that agents make to a central bank in the beginning of the time horizon, which then rewards them over time for their initial contribution. Moreover, the fact that they are imposed ex-ante is crucial; if this were not the case, then it would not be possible to guarantee everyone's mutually beneficial behavior over the whole time horizon without considering the sorts of "trigger" strategies discussed in Section 1.

To close this discussion, there are many ways policy makers may go about recouping the cost of policy implementation. Moreover, when the kinds of target paths that a policy maker is interested in inducing represent improvements over an inefficient status quo, a suitably chosen ex-ante cost sharing scheme ensures voluntary participation in the mechanism and ex-post individual rationality. At the same time, while the above comments lend the proposed policy instruments a reasonable degree of credibility, they come with the great caveat of ex-ante payments. For,

even though, under this policy, agents achieve ex-post Pareto gains over the status quo, the fact that they must incur all the (discounted) costs upfront might be a problem, especially if the time horizon is very long and the discount factor very small. But, even in that case, one would expect that the gains to be had are quite large, and so worth the ex-ante costs. We now briefly comment on a few further implications of Theorem 1.

The Benevolence of Linear Costs. In the special case of linear $c_i(\cdot)$ and $K_i(\cdot)$, we need not introduce the negatives of these functions in the mechanism (2), since the game already has a linear-state structure. Moreover, along the lines just discussed, the mechanism charges can be modified to

$$\beta_i(t)(x(t) - \hat{x}(t)) \text{ and } \gamma_i(x(T) - \hat{x}(T)).$$

Supposing that agents' linear costs are given by $c_i x(t)$ and $K_i x(T)$ where $c_i, K_i \in \mathfrak{R}$, these actions have the sole effect of transforming the co-state variable expression given by (6) to

$$\lambda_i^{\beta_i, \gamma_i}(t) = (\gamma_i + K_i)e^{-\int_t^T (\delta - g(s))ds} + \int_t^T (\beta_i(s) + c_i)e^{-\int_t^s (\delta - g(\tau))d\tau} ds.$$

Thus, the analysis pursued in Theorem 1 is virtually unchanged, so that in Markov-perfect open-loop equilibrium, we obtain the desired target path \hat{q} without any exchange of money whatsoever. In particular, there is no need to consider ex-ante cost sharing agreements to ensure budget balance and individual rationality. This point is particularly important in the context of Dutta and Radner's climate change model [9, 10, 11], as they restrict their attention to linear state-dependent costs.

Multi-Dimensional State Space. Theorem 1 can be extended to deal with a multi-dimensional state space in a straightforward way.⁹ In particular, granting multi-dimensional differentiability requirements, Assumption 1 and Assumption 2 can be relaxed to hold for at least one coordinate of the state vector and the mechanism charges for agent i , $\beta_i(\cdot)$ and γ_i , would be introduced to that coordinate only. Then the reasoning would proceed along similar lines. Of course, for an m -dimensional state space, the cancellation of the original cost functions $c_i(x_1, x_2, \dots, x_m)$ and $K_i(x_1, x_2, \dots, x_n)$ must be maintained in the mechanism's specification.

Interior Target Paths and Mechanism Uniqueness. If agent i 's target path \hat{q}_i is interior, then Assumption 2 applied to agent i can be relaxed to hold only for control-state paths in which

⁹For a concrete example see Athanassoglou et al. [1].

q_i is interior, and Theorem 1 would still hold. Moreover, in this case ensuring that $\beta_i(\cdot)$ and γ_i satisfy integral equation (8) for agent i is both necessary and sufficient for the two paths' restrictions to agent i to coincide. As a consequence, the restriction of the mechanism (β, γ) satisfying the statement of the theorem to agent i is unique.

In closing this section, we note that similar analysis as the one presented in Theorem 1 may be applicable to other games that admit Markov-perfect open-loop equilibria (see Clemhout and Wan [5], Reinganum [19], Dockner et al. [6], Fershtman [12] for a description and axiomatic characterization of such games).

5 An Application to Global Climate Change

5.1 A Differential Game Model of Global Climate Change

In this section we follow the model of Dutta and Radner [9, 10, 11] and apply the insights of Theorem 1 to a differential game of global climate change. To this end, suppose we have n countries who are involved in productive economic activity, which requires the emission of pollutants into the atmosphere. Country i 's emissions at time t are denoted by $e_i(t)$ and her resultant payoff by $h_i(e_i(t))$, where the function $h_i(\cdot)$ is strictly concave. The amount of pollution in the atmosphere at time t is denoted by $x(t) \in \mathfrak{R}_+$. Departing from Dutta and Radner, we make the assumption that country i incurs an arbitrary (i.e., not necessarily linear) cost that depends on the amount of pollution that is present in the atmosphere at that time. Thus, her net payoff at time t is given by

$$h_i(e_i(t)) - c_i(x(t)), \quad (13)$$

where $c_i(\cdot)$ is an arbitrary non-negative function. The global stock of pollution at time t , $x(t)$, is governed by the following law of motion

$$\dot{x}(t) = \sum_{i=1}^n e_i(t) - \sigma x(t), \quad (14)$$

where $\sigma > 0$ is the coefficient of natural atmospheric purification. The countries thus play the following differential game

$$\begin{aligned} \max_{e_i} \quad & \int_0^T e^{-\delta t} [h_i(e_i(t)) - c_i(x(t))] dt \\ \text{subject to:} \quad & \dot{x}(t) = e_i(t) + \sum_{j \neq i} e_j^*(t) - \sigma x(t) \\ & e_i \geq 0, \quad x(0) = \hat{x}. \end{aligned} \quad (15)$$

This is a simple special case of the class of games analyzed in this paper. It has a one-dimensional state space whose motion depends linearly on the controls and state (thus, it satisfies part (a) of Assumption 1), as well as an objective functional that is strictly concave in the control and control-state separable. Hence, it is relatively straightforward to apply the reasoning of Theorem 1 and obtain a simple closed-form solution to the integral equation given by (8). To this end, denote the globally optimal emission rate path over $[0, T]$ by

$$e^{GO} = \{(e_1^{GO}(t), e_2^{GO}(t), \dots, e_n^{GO}(t)) : t \in [0, T]\}.$$

The path e^{GO} is taken to be feasible (i.e., no negative emissions rates are allowed) and continuously differentiable, and corresponds to a “desirable” emissions trajectory that is determined according to whatever criteria of efficiency and equity the world mutually agrees upon.¹⁰ In what follows, we exhibit a mechanism that induces exactly this path in Markov-perfect equilibrium.

Theorem 2 *Consider the climate change differential game (15) and let e^{GO} be a globally optimal emissions path. The linear-state mechanism $(\beta^{GO}, \gamma^{GO})$ such that for $i \in \{1, 2, \dots, n\}$*

$$\begin{aligned} \beta_i^{GO}(t) &= h_i''[e_i^{GO}(t)](e_i^{GO})'(t) - (\delta + \sigma)h_i'[e_i^{GO}(t)], \quad t \in [0, T], \\ \gamma_i^{GO} &= -h_i'[e_i^{GO}(T)] \end{aligned}$$

*uniquely induces e^{GO} in Markov-perfect equilibrium. A global-optimum inducing linear-state mechanism is unique if the globally optimal emissions path is interior.*¹¹

Proof. Consider a linear-state mechanism (β, γ) where

$$\phi_i(x(t), t) = \beta_i(t)x(t) + c_i(x(t)), \quad \phi_i(x(T)) = \gamma_i x(T), \quad i \in \{1, 2, \dots, n\},$$

which induces the following linear state differential game among the n countries

$$\begin{aligned} \max_{e_i(\cdot)} \quad & \int_0^T e^{-\delta t} [h_i(e_i(t)) + \beta_i(t)x(t)] dt + e^{-\delta T} \gamma_i x(T) \\ \text{subject to:} \quad & \dot{x}(t) = e_i(t) + \sum_{j \neq i} e_j^*(t) - \sigma x(t) \\ & e_i \geq 0, \quad x(0) = \hat{x}. \end{aligned} \tag{16}$$

¹⁰The Global Pareto Optimum described in Dutta and Radner [9, 11] would be a reasonable candidate in this regard.

¹¹Since a globally optimal emissions path involving zero emissions is deeply unrealistic, the uniqueness result has bite.

Agent i 's current-value Hamiltonian is given by the following expression

$$H_i(e_i, e_i^*, x, \lambda_i, t) = h_i(e_i) + \beta_i(t)x + \lambda_i \left(e_i + \sum_{j \neq i} e_j^* - \sigma x \right).$$

The differential equation and terminal condition for the co-state variable is the following:

$$\begin{aligned} \dot{\lambda}_i &= (\delta + \sigma)\lambda_i - \beta_i(t) \\ \lambda_i(T) &= \gamma_i \end{aligned} \tag{17}$$

Solving (17) yields

$$\lambda_i^{\beta_i, \gamma_i}(t) = \gamma_i e^{-(\delta + \sigma)(T-t)} + \int_t^T e^{(\delta + \sigma)(t-s)} \beta_i(s) ds. \tag{18}$$

On the other hand, the Hamiltonian-maximizing condition is

$$\begin{aligned} h'_i(e_i) + \lambda_i^{\beta_i, \gamma_i} &\leq 0 \\ e_i(h'_i(e_i) + \lambda_i^{\beta_i, \gamma_i}) &= 0, \quad e_i \geq 0. \end{aligned} \tag{19}$$

As the Hamiltonian H_i is strictly concave in (e_i, x) , conditions (19) are necessary and sufficient for a Markov-perfect open loop equilibrium. Since the function h'_i is one-to-one, to prove our result it is both necessary and sufficient to find $\beta(\cdot)$ and γ so that

$$h'_i[e_i^{GO}(t)] = h'_i[e_i^{\beta, \gamma}(t)], \quad \text{for all } i \in \{1, 2, \dots, n\} \text{ and } t \in [0, T]. \tag{20}$$

Given Equation (18) and the equilibrium conditions (19) we may focus on agents separately and try to find $\beta_i(t)$ and γ_i such that

$$-h'_i[e_i^{GO}(t)] = \gamma_i e^{-(\delta + \sigma)(T-t)} + \int_t^T e^{(\delta + \sigma)(t-s)} \beta_i(s) ds, \quad \text{for all } t \in [0, T]. \tag{21}$$

Clearly, if e_i^{GO} is interior, then Condition (21) is necessary and sufficient for the two paths to coincide (i.e., $\beta_i(\cdot)$ and γ_i cannot be such that the inequality in (19) is strictly satisfied). To satisfy Equation (21) for $t = T$, we must set

$$\gamma_i^{GO} = -h'_i[e_i^{GO}(T)]. \tag{22}$$

This fixes our choice of γ_i^{GO} . Integral equation (21) is straightforward and may be solved by inspection. Suppose $\beta_i^{GO}(s) = h''_i[e_i^{GO}(s)](e_i^{GO})'(s) - (\delta + \sigma)h'_i[e_i^{GO}(s)]$. Then the right-hand-

side of (21) obtains

$$\begin{aligned}
RHS &= -h'_i[e_i^{GO}(T)]e^{-(\delta+\sigma)(T-t)} + \int_t^T e^{(\delta+\sigma)(t-s)} (h''_i[e_i^{GO}(s)](e_i^{GO})'(s) - (\delta + \sigma)h'_i[e_i^{GO}(s)])ds \\
&= -h'_i[e_i^{GO}(T)]e^{-(\delta+\sigma)(T-t)} + e^{(\delta+\sigma)t} \int_t^T \frac{d}{ds} \left(h'_i[e_i^{GO}(s)]e^{-(\delta+\sigma)s} \right) ds \\
&= -h'_i[e_i^{GO}(T)]e^{-(\delta+\sigma)(T-t)} + e^{(\delta+\sigma)t} \left(h'_i[e_i^{GO}(T)]e^{-(\delta+\sigma)T} - h'_i[e_i^{GO}(t)]e^{-(\delta+\sigma)t} \right) \\
&= -h'_i[e_i^{GO}(t)],
\end{aligned}$$

as desired. Repeating this argument for all countries establishes that the linear-state mechanism $(\beta^{GO}, \gamma^{GO})$ where

$$\begin{aligned}
\beta_i^{GO}(t) &= h''_i[e_i^{GO}(t)](e_i^{GO})'(t) - (\delta + \sigma)h'_i[e_i^{GO}(t)], \quad t \in [0, T], \\
\gamma_i^{GO} &= -h'_i[e_i^{GO}(T)]
\end{aligned} \tag{23}$$

for $i \in \{1, 2, \dots, n\}$ induces e^{GO} in open-loop, and therefore by Proposition 1, Markov-perfect equilibrium. ■

Comments. The mechanism prescribed by Theorem 2 depends solely on agents' individual payoff and cost functions and target emission paths. The fact that it is not explicitly influenced by their peers' analogous features is important, as it compels agents to focus on parameters that they have direct knowledge of and control over. Thus, one might expect that this scheme stands a better chance of being well implemented. Moreover, assuming that the control-state path (e^{GO}, x^{GO}) represents a Pareto improvement over an inefficient tragedy-of-the-commons status quo, the mechanism may be implemented with an ex-ante cost-sharing allocation to satisfy budget balance and individual rationality ex-post, along the lines discussed in Section 4.2. On this score, perhaps the relevant banking entity for these ex-ante cost payments would be an organization like the World Bank or the U.N. Environmental Programme. We repeat, however, that if costs are taken to be linear in the global stock of emissions, we can simply introduce the modification of mechanism $(\beta^{GO}, \gamma^{GO})$ that was discussed in Section 4.2 to the original cost function, and this policy will induce the global optimum without the need for cost-sharing agreements and the exchange of any money whatsoever.

A complete game-theoretic treatment of climate change would incorporate endogenous technological change and population growth into the model. With regard to technological change, the model we study could exogenously take it into account by making the payoff functions h_i depend

explicitly on time in addition to emissions. One would expect that an equal level of output could be generated with lower emissions as time increases and this could be modeled directly into the payoff function. Of course, such an approach disregards the endogenous costs of adopting and developing cleaner technologies. A rigorous treatment would involve modeling population growth and technological change into the objective function and state equations explicitly along the lines of Dutta and Radner’s early work [9, 10], and coupling it with a policy design framework.

6 Conclusion and Future Work

In this paper we have investigated a useful class of differential games from a “mechanism-design” point of view. Our main result establishes that, within this class, simple mechanisms which take into account only the state variable wield considerable power in credibly influencing equilibrium outcomes. An important property of the stipulated mechanisms is that they typically require moderate monitoring capabilities. Moreover, they can be implemented to satisfy budget balance and, where applicable, ex-post individual rationality. From a practical standpoint, our proposed economic instruments often result in simple closed-form policy recommendations, lending themselves to insightful predictive analysis. The proof technique we introduce uses the theory of integral equations and, we believe, may be of independent interest in studying mechanism design issues of tractable differential games. We demonstrate its applicability in a simple game of climate change to determine a mechanism that uniquely induces the globally optimal emissions path.

Interesting theoretical extensions of our work involve relaxations of our assumptions. As a first step, one may try a similar approach as the one presented in this paper for other differential games which admit Markov-perfect open loop equilibria. Or, more ambitiously, one may attempt to introduce (potentially high-order) control-state nonlinearity and non-separability to accommodate modeling assumptions that are natural in many important dynamic games. While remaining cautiously optimistic about our approach’s potential, we anticipate that this would complicate the analysis considerably. On a different note altogether, a fruitful area for further research would be to examine how the cost-sharing schemes proposed in Section 4.2 could be credibly implemented on some sort of dynamic, instead of ex-ante, basis.

With regard to the climate change application, incorporating stochastic uncertainty and catastrophic risk into the basic model along the lines of Xepapadeas [22] and Clarke and Reed [4] would be extremely interesting. On the applied front, one can make the climate change analysis concrete by using real data to determine the mechanism described in Theorem 2. For this purpose, the

calibrated numerical examples found in Dutta and Radner [9, 11] could be an immediate starting point.

References

- [1] S. Athanassoglou, W. T. Huh, G. Sheriff, and T. Siegfried (2009), “Simple Mechanisms for Groundwater Extraction Under Spatial Complexity and Economic Heterogeneity,” Columbia University, working paper.
- [2] S. Barrett (2003), *Environment and Statecraft: The Strategy of Environmental Treaty-Making*, Oxford University Press.
- [3] T. A. Burton (2005), *Volterra Integral and Differential Equations*, Elsevier, 2nd Edition.
- [4] H. Clarke and W. Reed (1994), “Consumption/Pollution Tradeoffs in an Environment Vulnerable to Pollution-related Catastrophic Collapse,” *Journal of Economic Dynamics and Control*, 18, 991-1010.
- [5] S. Clemhout and H. Y. Wan, Jr. (1974), “A Class of Trilinear Differential Games,” *Journal of Optimization Theory and Applications*, 14, 419-424.
- [6] E. Dockner, G. Feichtinger, and S. Jorgensen (1985), “Tractable Classes of Nonzero-Sum Open-Loop Nash Differential Games: Theory and Examples,” *Journal of Optimization Theory and Applications*, 45, 179-197.
- [7] E. Dockner, S. Jorgensen, N. Van Long, and G. Sorger (2000), *Differential Games in Economics and Management Science*, Cambridge University Press.
- [8] E. Dockner, N. Van Long, and G. Sorger (1996), “Analysis of Nash Equilibria in a Class of Capital Accumulation Games,” *Journal of Economic Dynamics and Control*, 20, 1209-1235.
- [9] P. K. Dutta and R. Radner (2004), “Self-Enforcing Climate-Change Treaties,” *Proceedings of the National Academy of Sciences*, USA, 101-14, 5174-5179.
- [10] P. K. Dutta and R. Radner (2006), “Population Growth and Technological Change in a Global Warming Model,” *Economic Theory*, 29, 251-270.
- [11] P. K. Dutta and R. Radner (2007), “A Strategic Analysis of Global Climate Change: Theory and Some Numbers,” *Journal of Economic Behavior and Organization*, forthcoming.

- [12] C. Fershtman (1987), "Identification of Classes of Differential Games for Which the Open-Loop is a Degenerate Feedback Nash Equilibrium," *Journal of Optimization Theory and Applications*, 55, 217-231.
- [13] C. Fershtman and S. Nitzan (1991), "Dynamic Voluntary Provision of Public Goods," *European Economic Review*, 35, 1057-1067.
- [14] S. Jorgensen, G. Martin-Herran, and G. Zaccour (2003), "Agreeability and Time-Consistency in Linear-State Differential Games," *Journal of Optimization Theory and Applications*, 119, 49-63.
- [15] S. Jorgensen and G. Zaccour (2001), "Time-Consistent Side Payments in a Dynamic Game of Downstream Control," *Journal of Economic Dynamics and Control*, 25, 1973-1987.
- [16] G. Martin-Herran and G. Zaccour (2005), "Credibility of Incentive Equilibrium Strategies in Linear-State Differential Games," *Journal of Optimization Theory and Applications*, 162, 367-389.
- [17] J. Reinganum (1981), "Dynamic Games of Innovation," *Journal of Economic Theory*, 25, 21-41.
- [18] J. Reinganum (1982), "A Dynamic Game of R and D: Patent Protection and Competitive Behavior," *Econometrica*, 50, 671-688.
- [19] J. Reinganum (1982), "A Class of Differential Games for Which the Closed-Loop and Open-Loop Equilibria Coincide," *Journal of Optimization Theory and Applications*, 36, 253-262.
- [20] S. P. Sethi and G. Thompson (2000), *Optimal Control Theory: Applications to Management Science and Economics*, 2nd Edition, Springer.
- [21] F. Wirl (1996), "Dynamic Voluntary Provision of Public Goods: Extension to Nonlinear Strategies," *European Journal of Political Economy*, 12, 555-560.
- [22] A.P. Xepapadeas (1992), "Environmental Policy and Dynamic Nonpoint-Source Pollution," *Journal of Environmental Economics and Management*, 23, 22-39.