Review of Economic Studies (2020xx) **0**, 1–0 doi:10.1093/restud/rdaa000

Markov-Perfect Equilibria in Differential Games — with an Application to Climate Policy

NIKO JAAKKOLA

University of Bologna, Piazza Scaravilli 2, 40126 Bologna, Italy; CESifo

FLORIAN WAGENER

CeNDEF, Amsterdam School of Economics, University of Amsterdam, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands; Tinbergen Institute; CESifo

E-mail: f.o.o.wagener@uva.nl

We analyse discontinuous Markovian strategies for differential games. The best response correspondence uniquely maps almost all profiles of opponents' strategies back to the strategy space. We thus make Markov-perfect equilibria in a wide class of differential games well-behaved, resolving a long-standing open problem. We provide a readily applicable necessary and sufficient condition for best responses and Markov-perfect Nash equilibria. We demonstrate our methods in a canonical model of non-cooperative mitigation of climate change. Our approach provides novel, economically important results: we obtain the entire set of symmetric Markov-perfect equilibria, and demonstrate that the best equilibria can yield a major welfare improvement over the equilibrium which previous literature has focused on. International climate negotiations can be seen as being about coordination on good equilibria, rather than about bargaining over the limited surplus available in a dynamic prisoner's dilemma.

1. INTRODUCTION

Climate change is widely recognised as a severe threat to economic and human wellbeing. Solving this problem has proven to be difficult: a stable climate is a global public good. Investment into such a public good, that is, refraining from burning fossil fuels, is costly, and all countries would prefer to free-ride on the efforts of others. Dynamic public good provision is typically seen to have the structure of a prisoner's dilemma, with disappointing outcomes.

However, dynamic games often have multiple Nash equilibria, of which some Pareto-dominate others. In this paper, we construct the entire set of symmetric equilibria in Markovian strategies to a class of dynamic public good games and apply the result to a classic game of dynamic climate mitigation (van der Ploeg and de Zeeuw, 1992). There is a single stock of a public bad—carbon in the atmosphere causing harmful climate change—and a set of symmetric players benefitting individually from emitting more carbon. Players choose Markovian strategies: policy rules which condition emissions only on the current carbon stock. There are many equilibria; coordinating on a good one is of first-order importance for welfare. Indeed, the strategic incentives embedded into policy rules can support non-cooperative, Markov-perfect equilibria which get close to, or even achieve, the socially optimal welfare outcome.[1] The welfare gains are substantial, relative to the equilibrium that has been—for reasons of tractability—the focus of the literature. This implies climate negotiations can leverage self-enforcing strategic incentives to achieve socially beneficial outcomes.

It has been argued that equilibrium multiplicity is important in public goods games (Dockner and Ngo Van Long, 1993; Dutta and Radner, 2004).[2] To the best of our knowledge, we are the first to make this argument in a continuous-time (i.e., differential game) framework which is game-theoretically sound. Our main innovation is to treat discontinuous Markovian strategies in a manner that makes the game well-specified. The continuous-time framework allows us to construct the entire set of symmetric Markov-perfect equilibria and discuss its welfare properties, showing that any continuous-value equilibrium Pareto-dominates the linear equilibrium. The latter has been the main focus of the literature. The best equilibria get remarkably close to the first-best.[3]

Our main contribution is methodological, and the main results two-fold: structural and constructive. Structurally, Theorem 3.1 shows that allowing discontinuous strategies makes the concept of Markov-perfect equilibrium (MPE) well-founded. The literature on differential games in economics has been unable to properly specify a strategy space such that payoffs can be computed for any strategy profile. Two approaches have been taken. One implies that best responses often fail to exist, as the “natural” best response does not belong to the specified strategy space. The other approach has been to allow a larger set of strategies, but rule out problematic strategy profiles, so that a player’s strategy set depends on the strategies simultaneously chosen by the other players (Bernhard, 2024). We put differential games on a solid game-theoretic foundation, solving, for the scalar state case, this open problem discussed in, for instance, the textbooks of Başar and Olsder (1982), Fudenberg and Tirole (1991), and Dockner et al. (2000)[4]. We expand on this below.

1. It is well-known that, under history-dependent strategies, good outcomes can be sustained by threats based on punishing undesirable past actions: our equilibria rely instead on punishments baked into the Markovian strategies (as in, e.g., Fudenberg and Tirole, 1983; Dutta and Sundaram, 1993).

2. Dutta and Sundaram (1993) demonstrate, in a discrete-time renewable resource game, a “good” outcome sustained by appropriate Markovian strategies, but do not give a full characterisation of the set of equilibria.

3. The continuous-time framework is essential to computing the entire set of (symmetric) equilibria. In discrete time, constructing the entire set—as opposed to presenting individual alternative equilibria—is to our knowledge difficult in a setting such as ours, because the optimal action depends on global properties of the value function.

4. Krasovskiĭ and Subbotin (1988) are an exception, as they do not place restrictions on strategies. They however work in the more restricted context of zero-sum games, and they consider the smaller class of constructive motions, rather than the larger class of generalised motions that we utilise.

Constructively, Theorem 3.2 provides a readily applicable necessary and sufficient condition for characterising the best response to a profile of opponents' strategies in a general class of public investment games. From this, MPEs can readily be constructed; we give a general condition for symmetric MPEs, Theorem 3.3, and demonstrate its use in the climate application. Our methods are relevant for many other applications, without the need to resort to special functional specifications, such as the commonly-used linear-quadratic framework.

Our climate policy application is an example of how our methods can be used to derive novel results for even well-understood models. We use the seminal model of non-cooperative mitigation of climate change of van der Ploeg and de Zeeuw (1992), with its linear-quadratic framework, so that a linear MPE exists. Dockner and Ngo Van Long (1993) proved the model also has non-linear "local" equilibria, but many of these could not be extended to the entire state space (as in Tsutsui and Mino, 1990). Rubio and Casino (2002) and Dockner and Wagener (2014) constructed even more local MPEs. However, it is not clear how policy rules undefined on the entire state space can be called equilibria (Rowat, 2007; Bernhard, 2024). Rowat (2007) enumerated all continuous, globally defined equilibria of this model. Schumacher et al. (2022), building on our approach, derive a complementary result (on which more below). In a renewable resource game, Sorger (1998) constructs a family of equilibria with a single discontinuity.[5]

We derive three main results for the application. First, we characterise the set of symmetric MPEs. The set is large, but construction of any member is straightforward (Theorems 4.7 and 4.8). Most of previously obtained equilibria are nested within the larger set of all symmetric equilibria, extending locally-defined equilibria into globally-defined, but discontinuous, MPEs.[6]

Second, we show that the new equilibria we identify are important. We focus on symmetric equilibria, which can be Pareto-ranked. The linear equilibrium—a central focus of the extant literature—is weakly Pareto-dominated by any other equilibrium with continuous value, regardless of the initial state. It also is dominated by discontinuous value equilibria if the initial state is small, but discontinuous value equilibria can be worse for large initial states. That is, the linear equilibrium yields the lowest surplus available in any MPE, except for equilibria featuring bad coordination failures.[7] The linear equilibrium involves adverse strategic incentives, as players try to exploit the fact that their opponents will reduce emissions in response to a higher carbon stock: emissions are everywhere dynamic strategic substitutes to each other. MPE with better welfare properties leverage dynamic strategic complementarity more to achieve lower long-run emissions. In a calibrated example, choosing the best equilibrium can close

5. Sorger (1998) is our closest precursor. We use a more general approach, and show that it is game-theoretically well-defined.

6. In our parameterisation, only one continuous globally defined equilibrium exists.

7. In a working paper version of this paper, we show that the same property holds if the utility function is isoelastic instead of linear-quadratic, as long as the carbon demand elasticity is sufficiently small. The qualitative properties of the discontinuous equilibria change at the same elasticity threshold as those of the continuous equilibrium (Wirl, 2014).

between 50–100% of the welfare difference between the linear equilibrium and the first-best outcome.[8] The question of which equilibrium is played is therefore of first-order importance for welfare.

The best equilibria achieve high welfare by a mechanism similar to trigger strategies: players are incentivised to adhere to collectively benign actions—namely, lowering their emission levels to "net zero" when the state reaches a low-carbon steady state—by an implicit threat: if any player were to emit more, so that the carbon stock would move above the benign steady state, all players would increase their emissions in response (as prescribed by the equilibrium strategy, and in mutual best response to each other). The economy would drift off to a higher long-term carbon stock. The threat of this makes the low emissions a best response.

We emphasise that, as MPE strategies are not conditioned on past actions, the trigger is defined in terms of the stock. The players thus coordinate on an equilibrium strategy, which can be argued to represent a self-enforcing agreement (Dockner and Ngo Van Long, 1993). There are many MPE; in ones which lead to particularly good welfare outcomes, the punishments are baked into the Markovian policy rules, which of course remain best responses to each other from any state. Thus coordination on a good equilibrium is very important. Markovian trigger strategies of the kind we discuss were first presented in Fudenberg and Tirole (1983). Similar strategies have been used, in a discrete-time context, also in environmental and resource economic applications (Dutta and Sundaram, 1993).[9] These applications discuss individual equilibria; we show how to construct the set of all equilibria, including ones without such triggers.

In a benign steady state of a constrained-efficient equilibrium, the players respond differently if the state falls slightly below the steady state level, compared to the state rising slightly above the steady state level. A fall in the carbon stock below the steady state leads to a discontinuous increase in emission rates, so that any spare "capacity" is rapidly eaten up by the players, returning the economy to the steady state. If the carbon stock increases to above the steady state, this results in a small increase in emission rates, causing the economy to drift to higher carbon concentrations and into an eventual more highly polluted steady state. The threat of this adverse outcome ensures that players prefer to remain in the benign steady state.

Third, we show that there exist equilibria involving extreme coordination failures; at high levels of atmospheric carbon, these can yield welfare outcomes even worse than the linear equilibrium. The coordination failure involves discontinuities in not just the actions, that is, the emission rates, but also in the value function. Such a discontinuity is only possible if no player can single-handedly move the economy from the high-carbon, low-value side of the discontinuity to the low-carbon, high-value side. Think, for example, of a large number of relatively small symmetric countries: if other countries are emitting heavily, then no single country can stop the rise of atmospheric carbon even if it was to stop

8. Schumacher et al. (2022) show that an equilibrium with a value function arbitrarily close to the value envelope, i.e. yielding the maximal value for any initial state achievable with any MPE, can be constructed. However, this requires that the number of discontinuities becomes arbitrarily high.

9. The threat is only credible if no player can unilaterally move the stock back to the benign steady state. Fudenberg and Tirole (1983) assume investment is irreversible; in Dutta and Sundaram (1993), the discrete nature of time ensures that unilateral action cannot reverse the state after a small deviating action. In our continuous-time model, near the benign steady states the each player could, but does not want to, undo the emissions which have led to the Markovian threat to be triggered.

emitting completely. In equilibrium it is then a best response to also emit heavily. All players, following the same logic, emit heavily only as an optimal response to a collective coordination failure.[10]

Before turning to the model, we discuss the technical complication which has not been addressed by the previous literature, why it has presented problems, and how we solve these. The complication relates to what kinds of policy functions should be allowed. In particular, best-response policy functions can generically be discontinuous in the state; but discontinuous policies can lead to nonexistence of classical solutions to the differential equation governing state evolution. This causes a difficulty: without a state trajectory, payoffs cannot be computed.

The literature has taken two approaches to this. The first is to allow only continuous Markovian strategies, ensuring existence and uniqueness of trajectories (e.g., Dockner and Ngo Van Long, 1993; Wirl, 1996, 2014). Two problems ensue: first, it leads to non-existence of best responses, in the space of continuous functions, to many strategy profiles of other players, as the "natural" best response would be discontinuous.[11] Second, allowing only continuous strategies may rule out, a priori, equilibria which are technically unproblematic. In our application, this matters for welfare: any discontinuous strategy equilibrium with a continuous value function Pareto-dominates the unique continuous equilibrium.

The second approach (e.g., Sorger, 1998; Dockner et al., 2000; Rowat, 2007) allows discontinuous strategies in principle, but rules out strategy *profiles* leading to technical complications. A restriction to "admissible" strategy profiles implies that a player's allowable set of strategies depends on the strategies simultaneously chosen by the other players: the model is then no longer a noncooperative game (in the sense of Nash, 1951).[12]

In the present paper, we extend the definition of a differential game to allow for strategies with discontinuities. Under our definition, state trajectories and payoffs can always be computed: they are so-called Filippov solutions to the ordinary differential equation governing the state dynamics (Filippov, 1988). We can thus allow strategy profiles to simply be product sets of the individual strategy spaces, as is standard in non-cooperative game theory. Our first main result states that each player has a unique, Markovian best response to almost every strategy profile of the other players (Theorem 3.1). These two features jointly solve the above, long-standing technical problem with MPEs in

10. This is of course not a renegotiation-proof outcome. In fact, all of our equilibria feature some initial states from which the equilibrium trajectories are not renegotiation-proof were it possible to pause the evolution of the dynamics and play again. We discuss this further in Section 5.

11. The common focus on symmetric equilibria hides this problem. However, in some games, the continuity requirement implies the non-existence of any globally defined symmetric equilibria (Kossioris et al., 2008).

12. This problem has recently been pointed out by Bernhard (2024). Klein and Rady (2011) discuss admissibility extensively. In their model, equilibria are sometimes sustained by the technical requirement of admissibility. Schumacher et al. (2022) provide a verification theorem—a sufficient condition—for a symmetric Markovian policy rule to constitute an MPE (essentially the same approach as in the previous literature). However, they do not cleanly specify an individual player's problem, and their approach does not determine payoffs for "problematic" strategy profiles. Thus their problem cannot truly be said to be a non-cooperative game. Their approach also cannot answer questions about, for instance, the set of equilibria.

differential games for the single state variable case, by guaranteeing that payoffs can always be computed and best responses exist almost always.[13]

Our second main result, Theorem 3.2, gives a readily applicable necessary and sufficient condition for characterising the best response to a profile of opponents' strategies. This condition supplements the maximum principle with local requirements that have to hold at a discontinuity. These additional requirements restrict which discontinuities are possible at a given state. Imposing a fixed point condition on the best response correspondence, a Nash equilibrium in Markov-perfect strategies can then be constructed. Theorem 3.3 formulates the resulting necessary and sufficient conditions for symmetric equilibrium profiles.

To obtain our results, we use the theory of viscosity solutions to Hamilton–Jacobi–Bellman equations (Bardi and Capuzzo-Dolcetta, 2008), building on and developing ideas by Barles et al. (2013, 2014).[14] Because our interests are economic and game-theoretic, we primarily focus on policies rather than values. In particular, our methods directly construct the Markovian policy function that is optimal in the face of discontinuous dynamics, where the discontinuities arise from the strategies of the other players; of course the results hold as well for discontinuities resulting from, for instance, regime shifts. Unlike Barles et al. (2013, 2014), we cannot assume controllability of the dynamics near discontinuities without placing undue restrictions on the admissible set of strategies of the other players. Finally, the application of our results to a fixed point of the best response correspondence, to construct conditions for MPEs, is entirely novel.

In our definition of a differential game, we need to specify how the dynamics are resolved, and payoffs computed, in situations at which the dynamics involve discontinuities. In particular, we need to define meaningful payoffs for "chattering" or "relaxed" actions, where a player's actions may switch infinitely often in time. However, while our specification must allow for such behaviour, our results show that such outcomes never occur in equilibrium. Indeed, all the MPE we identify have perfectly standard classical solutions to the state dynamics.

Our results apply to numerous differential games; our example suggests that our methods can yield novel insights into all of these. Our methods also relax the restrictions on parametric model specification needed for tractability, making it straightforward to study models with novel specifications. We discuss other possible extensions in Section 5.[15]

The paper proceeds as follows. Section 2 sets up the basic model. The main results, existence and characterisation of the best response, are given in Section 3. We show how these are used in an application to climate change in Section 4. Readers primarily interested in the transboundary pollution game can read the beginning of Section 2 and then skip to Section 4, referring to previous sections when necessary. Section 5 wraps up. Proofs of the main results are relegated to the Online Appendix.

13. We thus solve the critique of Bernhard (2024). A large fraction of the literature on differential games works in a single-state setting. Extension of our results to more dimensions is left for future work.

14. We also rely on the theory of nonlinear dynamical systems to show that the best response is Markovian.

15. Many learning models specify a learning process, sometimes requiring parametric restrictions, in which the dynamics only move in one direction between the times when new information arrives, specifically to avoid the problems which we tackle (Keller and Rady, 2015; Sun, 2024); or the state dynamics are restricted by imposing an admissibility requirement on strategies (Klein and Rady, 2011) or beliefs (Board and Meyer-ter Vehn, 2013; Hauser, 2024). Our methods obviate the need for such requirements, at least for the class of games we consider.

2. THE MODEL

We consider MPEs in a differential game between N players. Player i's problem is to maximise

$$\int_0^\infty \exp(-\rho_i t) u_i(x(t), a_i(t)) dt \tag{2.1}$$

where $u_i(x(t),a_i(t))$ is flow utility, a function of the current state variable $x(t) \in \mathcal{X} = [x_{\min}, x_{\max}] \subset \mathbb{R}$ and player i's own current action $a_i(t) \in \mathcal{A}_i = [a_{\min,i}, a_{\max,i}] \subset \mathbb{R}$. The discount rate $\rho_i > 0$ is constant. The maximisation is conducted over trajectory–action pairs (x, a_i), with $x: [0,\infty) \to \mathcal{X}$ absolutely continuous and $a_i : [0,\infty) \to \mathcal{A}_i$ measurable, satisfying

$$\dot{x}(t) = a_i(t) + g_i(x(t)), \qquad x(0) = x_0, \tag{2.2}$$

for almost all $t \geq 0$. Here $g_i(x) = \Phi_{-i}(x) + h(x)$ is the sum of the physical dynamics $h(x)$ as well as the aggregate action $\Phi_{-i}(x) \equiv \sum_{j \neq i} \phi_j(x)$ of the other players. The functions ϕ_j are Markovian policy rules: player $j \neq i$ sets, at all times, its action according to $a_j(t) = \phi_j(x(t))$. Player i's problem depends only on the aggregate action Φ_{-i}.[16]

Pollution build-up models are our main focus: the state x is a pollution stock, utility $u_i(x, a_i)$ is decreasing in the stock, that is $\frac{\partial u_i}{\partial x} \leq 0$, and actions a_i are emissions. The model is easily adapted to pollution abatement, accumulation of joint capital, and contributions to public goods; these adaptations entail only changes of signs.

Our primary interest is in MPEs, which requires us to construct player i's optimal policy, its best response to the other players' strategies, as a feedback (or Markovian) policy rule so that $a_i(t) = \phi_i(x(t))$ for all $t \geq 0$.[17] An MPE is a strategy profile $\phi = (\phi_1, \ldots, \phi_N)$ such that each player's strategy is a best response to the other players' strategies, from any initial state. For the notion of MPE to be well-defined, we have to be able to compute the payoffs for any strategy profile. Furthermore, the strategy space, the class of functions which the ϕ_i are allowed to belong to, should be *strategically complete*: for every player i and every profile $\phi_{-i} = (\phi_1, \ldots, \phi_{i-1}, \phi_{i+1}, \ldots, \phi_N)$, the optimal action schedule a_i^* should be generated by a strategy ϕ_i from the strategy space.[18] That is, we request the strategy space to be sufficiently large to accommodate all best response action schedules. This is a consistency requirement, particular to differential games, connecting the action schedule space to the strategy space.

The literature has not addressed these desiderata: nearly each model we know of fails to satisfy at least one.[19] Our aim is to meet both. To do so, in the remainder of this

16. In particular, player i's best response is identical to any set of opponents' strategies which involve the same aggregate action.

17. Our results of course also apply to a single-agent optimal control problem under discontinuous dynamics.

18. We focus on autonomous feedback strategies: strategies conditioned only on the state variable, not calendar date. We do not discuss open-loop strategies, that is, strategies which give actions as functions of time. It is well-known that open-loop strategies are not robust to unforeseen one-off shocks to the state: the resulting Nash equilibria lack strong time consistency, a concept similar to subgame perfection (Dockner et al., 2000).

19. The locally-defined MPEs of Dockner and Ngo Van Long (1993) do not allow computation of payoffs from certain initial states. Discontinuous strategies with an admissibility restriction on strategy

section we allow strategies that are discontinuous in the state. In Section 3 we show that, after the extension, best replies almost always exist; they can be fully characterised; and consequently symmetric MPEs can be characterised in a manner which allows us to construct the set of all equilibria, as demonstrated for an application in Section 4.

2.1. *Strategies*

In the remainder of this section, we consider the problem of a single player, and we therefore drop the indices i and $-i$ until further notice. We want to construct a Markovian best response ϕ to an aggregate action Φ. The key question is the space of functions to which the strategy ϕ belongs: we discuss why a natural function space must admit discontinuous functions.

The *Hamiltonian* of player i is

$$H(x,p)=\max_{a\in\mathcal{A}}[u(x,a)+p(a+g(x))]. \tag{2.3}$$

This maximisation problem is the static problem which player i must solve at each point in time; jumping slightly ahead, when player i's value function V is differentiable, the maximisation is conducted under $p=V'(x)$.

The first-order condition for an interior maximum to player i's static problem is

$$\frac{\partial u}{\partial a}(x,a)+p=0.$$

To consider only situations with unique maximisers, we assume $\frac{\partial^2 u}{\partial a^2}<0$ everywhere.[20] We can then solve for the optimal action in terms of state and costate as

$$a=a^*(x,p). \tag{2.4}$$

However, the dependence of the optimal policy $\phi(x)=a^*(x,V'(x))$ on the state is neither smooth nor necessarily continuous. First, because the action space is compact, we have corner solutions to the maximisation in (2.3), with the action passing from interior values of $\mathcal{A}$ to boundary, or corner, values when the costate passes a *switching point* (we italicise terms which we use later in the paper). This implies that the allowable strategy $\phi(x)$ should admit kinks.

Second, it is well-known that if the dynamics are continuous and smooth but non-convex, there may exist a state $\hat{x}$ such that the maximisation problem in Equation (2.3) has two distinct solutions, for values of the costate p equal to the left and right limits of $V'(x)$ at $\hat{x}$. We call $\hat{x}$ a Skiba, or *indifference*, state (Skiba, 1978; Wagener, 2003). At such a point, there are two optimal solutions—one moving towards lower values of the state, the other towards higher values, both yielding the same value at $\hat{x}$—so that the optimal

profiles (Dockner et al., 2000) implies available strategies depend on other player's chosen strategies. And a requirement of continuity means the best response often does not exist in the space of continuous policies.

20. That is, we do not consider discontinuities due to flow utility not being strictly concave in the action.

Markovian policy ϕ is discontinuous. Such situations happen for a broad range of the opponents' aggregate action Φ, even if that action is smooth. In other words, the optimal policy response to smooth and continuous Markovian strategies cannot, in general, be described with a continuous Markovian strategy.

The extant literature either assumes *a priori*, often implicitly, that strategies must be continuous, in which case best responses to other players' strategies do not exist if player i's optimal policy features a Skiba point— and this occurs generically; or it imposes an admissibility or "feasibility" restriction which effectively conditions player i's strategy set on strategies simultaneously chosen by the other players. Our framework does not need these highly restrictive assumptions.

We allow a strategy ϕ, or any of its higher order derivatives, to have a finite number of discontinuities, which we call *jump points*.[21] Further, we require strategies to be piecewise *real analytic*, assuming that they can be expanded at each non-jump point in a power series with a positive radius of convergence. This implies strategies are infinitely differentiable outside of a jump point.[22] Finally, at jump points we require that strategies and their first order derivatives have finite one-sided limits when approaching the jumps. Formally, we introduce the strategy space $\mathscr{S}$ as follows.

Definition 2.1. *The strategy space $\mathscr{S}$ consists of functions $\phi : \mathcal{X} \to \mathcal{A}$ that are real analytic, except at a finite number of points, and have one-sided derivatives everywhere.*

2.2. *Dynamics and payoffs*

Allowing discontinuous strategies requires a reinterpretation of the evolution equation (2.2). Supposing player i chooses a Markovian policy, the equation becomes $\dot{x}(t) = \phi(x(t)) + g(x(t))$. As long as the right hand side of this equation is a Lipschitz-continuous function of the state, that is, continuous and almost everywhere differentiable with a uniformly bounded derivative, the Picard-Lindelöf theorem ensures the existence and uniqueness of a state trajectory x solving the equation. This is why the existing literature has often required the strategies ϕ to be Lipschitz-continuous: it guarantees payoffs can be computed. With a discontinuous right hand side, neither existence nor uniqueness of classical solutions can be guaranteed, and payoffs cannot be determined for all strategy profiles.

The main problem can be illustrated by the simple example

$$\dot{x}(t) = f(x(t)), \qquad f(x) = \begin{cases} f_-, & x < 0, \\ f_+, & x \geq 0, \end{cases} \tag{2.5}$$

where the right hand side f has a jump discontinuity at $x = 0$. Consider the special situation that $f_- > 0$ and $f_+ < 0$. Solutions to this equation increase at a constant rate

21. The mathematics for infinitely many jumps would be substantially more complicated, and we see policy functions with infinitely many jump points as having less practical relevance (but see Ivanov, 1990).

22. If the state–costate system of the optimisation has a steady state, and the function g is k times differentiable, then the best response policy function can only be guaranteed to be $k-1$ times differentiable. To map the best response back to the same function space, infinite differentiability is thus required. Real analyticity is a technical assumption and not particularly restrictive.

as long as $x(t)$ is negative, and decrease at a constant rate if $x(t)$ is positive. Using the intuitive terminology of Barles et al. (2013), we call this a *push-push* situation, as the state is pushed towards $x=0$ by the dynamics. It is natural to expect that if the state arrives in $x=0$ at some point in time, it will continue to remain at the origin for all future times. But $x(t)=0$ is not a solution to the differential equation. In our context f_+ and f_- would represent the limits of the dynamics arising from all players, including player i, using a discontinuous Markovian policy. The solution to this technical problem matters: player i's flow payoff at the steady state $x(t)=0$ depends on what is going on, in terms of player i's action, at that point.

For a heuristic explanation of how we resolve this issue, set $x(0)=0$, and suppose that a decision maker sets $\dot{x}(t)$ equal to either f_- or f_+. In particular, assume that if $\dot{x}(t_0)=f_-$, then this rate is maintained until $x(t)$ attains the level $x=\delta$, at which point $\dot{x}(t)$ is switched to f_+; here $\delta>0$ is some tolerance parameter. Likewise, if $\dot{x}(t_0)=f_+$, this rate is kept until $x(t)$ crosses the level $x=-\delta$ and then it is switched to f_-.

The state $x(t)$ is captured in an interval, bouncing between the boundary points $x=-\delta$ and $x=\delta$. During a fraction μ_- of the time, the state increases at a rate f_-; during a complementary fraction μ_+, it decreases at a rate f_+. These time fractions are inversely proportional to the absolute values of the rates, which implies

$$\mu_- = \frac{\frac{1}{|f_-|}}{\frac{1}{|f_-|}+\frac{1}{|f_+|}} = \frac{|f_+|}{|f_+|+|f_-|}, \qquad \mu_+ = 1-\mu_-. \tag{2.6}$$

For all tolerances $\delta>0$, player i's flow payoff is a weighted sum of the left action payoff, obtained when $\dot{x}(t)>0$, and right action payoff, obtained when $\dot{x}(t)<0$, with the weights given by the fraction of time the trajectory spends using these actions, respectively, μ_- and μ_+.

Taking the limit $\delta\to 0$, the state trajectory converges uniformly to the natural solution $x(t)=0$ for all t. This solution is formalised mathematically by replacing the differential equation (2.5) by a differential inclusion $\dot{x}(t)\in F(x(t))$. For this we introduce a set-valued function F such that $F(x)=\{f(x)\}$ if $x\neq 0$, and we extend it to $x=0$ such that $F(0)$ is convex and the graph of F is closed, that is, $F(0)=[f_-,f_+]$, where the interval $[f_-,f_+]$ connecting f_- and f_+ is defined as $\{\theta f_- + (1-\theta)f_+ : 0\leq\theta\leq 1\}$, irrespective of whether $f_- < f_+$ or not. The differential equation is then replaced by the differential inclusion for almost all t. This specification allows for the solution $x(t)=0$. These generalised solutions are known as "chattering", "sliding" or Filippov solutions. [23]

If $f_+>0$ and $f_-<0$ in (2.5), we have a *pull-pull* situation. The differential inclusion $\dot{x}(t)\in F(x(t))$, with initial condition $x(0)=0$, has now multiple solutions: $x(t)=f_+t$, $x(t)=f_-t$, and $x(t)=0$ for $0<t<t_0$ and $x(t)=f_\pm t - f_\pm t_0$ for $t\geq t_0$ are all different solution trajectories. The first two are *regular* solutions, while the others are *singular*: the latter solutions are not stable under small variations of the initial state. Multiple regular solutions typically occur as optimal trajectories originating at an indifference point, and so are perfectly standard.

In general, replacing differential equations with discontinuous right hand sides by differential inclusions that convexify the dynamics at discontinuities ensures the existence

23. Because we work in one dimension, these solutions are sliding in a degenerate sense.

of solutions (Filippov, 1988). Uniqueness is not guaranteed, but this is less problematic in our context.

Extending differential inclusions to controlled systems is straightforward, as the controlled evolution equation (2.2) itself can be naturally reformulated as a differential inclusion

$$\dot{x}(t) \in F(x(t)), \qquad x(0) = x_0. \tag{2.7}$$

First, for all points x where g is continuous, we set $F(x) = \mathcal{A} + g(x) = \{a + g(x) : a \in \mathcal{A}\}$. That is, the set $F(x)$ contains the rates of state increase available to player i at point x. We then extend F to any points of discontinuity by requiring that the graph of F is closed, that $F(x)$ is convex for all x, and that F is upper semi-continuous. This uniquely determines F.

To be specific, let x be a point of discontinuity of g, and let $g_-(x)$ and $g_+(x)$ be the left-hand and right-hand limits of $g(y)$ as $y \uparrow x$ or $y \downarrow x$ respectively. Then $F(x)$ is the convex hull of the union of the sets $F_-(x) = \mathcal{A} + g_-(x)$ and $F_+(x) = \mathcal{A} + g_+(x)$, and the elements $\dot{x} \in F(x)$ can be represented, with $a \in \mathcal{A}$ and $0 \le \mu \le 1$, as convex combinations [24]

$$\dot{x} = \mu(a + g_-(x)) + (1-\mu)(a + g_+(x)) = a + \mu g_-(x) + (1-\mu) g_+(x). \tag{2.8}$$

An *(admissible) trajectory* of the controlled system is an absolutely continuous function x for which (2.7) holds for almost all $t \ge 0$.

In the differential inclusion approach, trajectories are the primary objects. We need to determine the associated actions in order to evaluate the objective. By the Filippov selection theorem (Vinter, 2000, Theorem 2.3.13), if g has discontinuities $\bar{x}_1, \ldots, \bar{x}_J$, there are measurable functions a, $\mu_1, \ldots, \mu_J$, such that for almost all t

$$\dot{x}(t) = a(t) + g(x(t)), \tag{2.9}$$

if $x(t) \notin \{\bar{x}_1, \ldots, \bar{x}_J\}$, and

$$\dot{x}(t) = a(t) + \mu_j(t) g_-(x(t)) + (1 - \mu_j(t)) g_+(x(t)) \tag{2.10}$$

if $x(t) = \bar{x}_j$. The pair (x, a) is an *(admissible) trajectory–action pair* for the differential inclusion (2.7) if equations (2.9) and (2.10) hold for appropriately chosen weights.

To avoid technical complications, we consider differential games whose state space $\mathcal{X}$ is a compact interval $[x_{\min}, x_{\max}]$; we therefore have to specify what happens if the state leaves the state space. More precisely, we take it that the sum $g(x) = \Phi(x) + h(x)$ of aggregate action and system response is only defined for x in the interior $\mathring{\mathcal{X}}$ of $\mathcal{X}$. We extend the function g to all real numbers by setting $g(x) = g_+(x_{\min})$ if $x \le x_{\min}$ and $g(x) = g_-(x_{\max})$ if $x \ge x_{\max}$. This implies a corresponding extension of the set-valued function F. We redefine a trajectory to be admissible if $x_0 \in \mathcal{X}$, and if the differential

24. This is different from Barles et al. (2013, 2014), where the actions a_- and a_+ on both sides of a discontinuity take values in different spaces A_- and A_+.

inclusion (2.7) is satisfied almost everywhere. For an admissible trajectory, the *exit time* from $\mathcal{X}$ is[25]

$$\Theta = \inf\{t \geq 0 : x(t) \notin \mathcal{X}\}.$$

2.3. *Values*

Having obtained the trajectories and the actions, we can evaluate the objective. At the exit time, the state leaves the state space, the game stops, and the player receives a boundary payoff $\beta(x)$, which is equal to $\beta(x_{\min})$ if $x \leq x_{\min}$ and to $\beta(x_{\max})$ if $x \geq x_{\max}$. That is, instead of (2.1), we consider the objective

$$U(x,a) = \int_0^\Theta \exp(-\rho t) u(x(t),a(t)) dt + \exp(-\rho\Theta)\beta(x(\Theta)).$$

The payoff of the game for the player is

$$V(x_0) = \sup U(x,a),$$

where the supremum is taken over all admissible trajectory–action pairs (x,a) with initial state $x(0) = x_0$. The supremum is actually attained, and can thus be replaced by a maximum: the corresponding pair (x,a) is then an *optimal trajectory-action pair*. We note that $V(x) = \beta(x_{\min})$ if $x < x_{\min}$ and $V(x) = \beta(x_{\max})$ if $x > x_{\max}$.

2.4. *Feedback best response strategies*

From this point onwards, we reinstate the indices i and $-i$. We want to understand whether the optimal trajectory-action pair of player i—a time trajectory—can be generated by a Markovian policy $\phi_i \in \mathscr{S}$, for any initial state x_0, so that we can say that ϕ_i is a *Markovian best response* to $\phi_{-i} \in \mathscr{S}^{N-1}$. Existence of Markovian best responses consistent with all optimal trajectory-action pairs is what we term strategic completeness. We first need to define exactly what a Markovian best response is; the difficulty clearly lies at the points at which ϕ_{-i} is discontinuous.

We consider a strategy profile $\phi \in \mathscr{S}^N$. Because of the discontinuities, we again use a differential inclusion to describe the dynamics. We have to define how payoffs to a Markovian strategy are computed at points of discontinuity.

We set $f^\phi(x) = \phi_i(x) + \Phi_{-i}(x) + h(x)$ and introduce a set-valued map by $F^\phi(x) = \{f^\phi(x)\}$ for all points x where $f^\phi(x)$ is continuous. As before, we extend F^ϕ to any point of discontinuity by requiring that the graph of F^ϕ is closed, that $F^\phi(x)$ is convex for all x, and that F^ϕ is upper semi-continuous. For a point of discontinuity x, this implies that $F^\phi(x)$ is the closed interval $[f^\phi_-(x), f^\phi_+(x)]$ connecting $f^\phi_-(x)$ to $f^\phi_+(x)$. Moreover, the set-valued map is extended past the boundaries of $\mathcal{X}$ by setting it equal to its value at the nearest boundary point.

25. Thus we properly consider a free end time problem, rather than an infinite horizon problem (see the critique by Bernhard, 2024). In many problems an appropriate specification of the state space eliminates any meaningful difference between the two; see our application in Section 4.

If for almost all t we have

$$\dot{x}(t) \in F^{\phi}(x(t)), \qquad x(0) = x_0, \tag{2.11}$$

then x is an admissible Markovian trajectory when the player chooses the feedback strategy ϕ_i in response to ϕ_{-i}.

To compute payoffs, we need to specify utility. At points x where f^{ϕ} is continuous, utility is

$$u_i^{\phi}(x) = u_i(x, \phi_i(x)).$$

If f^{ϕ} is not continuous at x, then $\dot{x} \in F^{\phi}(x)$ can be written uniquely as a convex combination

$$\dot{x} = \mu(\phi_{i,-}(x) + g_{i,-}(x)) + (1-\mu)(\phi_{i,+}(x) + g_{i,+}(x)) \tag{2.12}$$

with $0 \le \mu \le 1$ again being the time fraction the solution spends on the left hand side of the discontinuity. Player i's utility is then

$$u_i^{\phi}(x) = \mu u_i(x, \phi_{i,-}(x)) + (1-\mu) u_i(x, \phi_{i,+}(x)). \tag{2.13}$$

With this choice, player i's payoff is

$$V_i^{\phi}(x_0) = \sup \left[\int_0^{\Theta} \exp(-\rho_i t) u_i^{\phi}(x(t)) dt + \exp(-\rho_i \Theta) \beta_i(x(\Theta)) \right], \tag{2.14}$$

where the supremum is taken over all solutions of (2.11), including singular solutions.

We emphasise that, under the differential inclusion approach, state and action trajectories, and thus values, can be computed for all strategy profiles ϕ, so that individual strategies can be freely chosen from $\mathscr{S}_i$, independently of the other players' strategies, in line with the notion of noncooperative games (Nash, 1951; Bernhard, 2024).

If, for a strategy profile ϕ, player i's payoff coincides with the payoff achieved by the optimal trajectory-action pair in response to ϕ_{-i}, i.e. $V_i^{\phi}(x_0) = V_i(x_0)$ for all x_0, we say that ϕ_i is a *Markovian best response* to ϕ_{-i}.

The key difference between equations (2.10) and (2.12) is that the latter requires a control to be a convex combination of the left and right limiting Markovian control. Then player i's payoff is invariant to pointwise deviations in ϕ_i, and a Markovian strategy ϕ_i encodes all the information required to compute payoffs, as the weight μ can be computed from the limits $\phi_{i,-}(\bar{x})$, $\phi_{i,+}(\bar{x})$ for any $\bar{x}$. We need to show that even under these restrictions Markovian best responses always exists; this is the first of our main results, to which we turn now.[26]

26. To see that this is not obvious and instead depends on further restrictions of the model, consider a symmetric game with $N=3$, discount rate $\rho=1$, utility $u(x,a_i) = a_i^2 - x^2$, dynamics $\dot{x} = \sum_{i=1}^{N} a_i$, state space $\mathcal{X} = \mathbb{R}$ and action space $\mathcal{A} = [-1/2, 1]$. Let $\bar{x} = 2/3$ and assume that $\phi_2 = \phi_3 = \phi$ with $\phi(x) = 1$ for $x < \bar{x}$ and $\phi(x) = -1/2$ for $x > \bar{x}$. The value function of player 1 satisfies $V_1'(x) > -1/2$ if $x < \bar{x}$ and $V_+'(x) < -1/2$

3. RESULTS

This section states the three main results of our article. The first asserts the well-definedness of the best-response correspondence for almost all strategy profiles. We first make precise our meaning of "almost all" by defining the concept of a "shy" set. Shy sets are analogues of measure zero sets in infinite-dimensional spaces, such as function spaces.

Let S be a subset of a complete metric linear space $\mathscr{V}$. The set S is *nowhere dense*, that is, "topologically small", if its complement contains an open and dense set; it is *shy*, or "measure-theoretically small", if there exists a Borel set B containing S and a measure on $\mathscr{V}$ that takes a finite value on some compact set, such that $\text{meas}(B-v)=0$ for all $v\in\mathscr{V}$: here $B-v=\{b-v:b\in B\}$. Note that if $\mathscr{V}$ is finite dimensional, a set is shy if and only if its Lebesgue measure is 0 (Hunt et al., 1992).

A *collection of jump points* is a finite discrete set $\mathcal{J}=\{\bar{x}_1,\ldots,\bar{x}_J\}$ with $\bar{x}_1=x_{\min}$ and $\bar{x}_J=x_{\max}$. We denote by $\mathscr{F}_{\mathcal{J}}$ the set of functions that are real analytic on $\mathcal{X}\backslash\mathcal{J}$, such that their derivatives have finite one-sided limits everywhere. The strategy space $\mathscr{S}_{\mathcal{J},i}$ of player i consists of those functions $\phi_i\in\mathscr{F}_{\mathcal{J}}$ such that $\phi_i(x)\in\mathcal{A}_i$ for all $x\in\mathcal{X}$. The strategy space $\mathscr{S}_{\mathcal{J},-i}\subset\mathscr{F}_{\mathcal{J}}^{N-1}$ of all players except player i is defined similarly. The strategy spaces $\mathscr{S}_i$ and $\mathscr{S}_{-i}$ are the union of, respectively, the sets $\mathscr{S}_{\mathcal{J},i}$ and $\mathscr{S}_{\mathcal{J},-i}$ over all collections $\mathcal{J}$. The spaces $\mathscr{F}_{\mathcal{J},i}$ and $\mathscr{F}_{\mathcal{J},-i}$ can be given the structure of a complete metric linear space, see e.g. Krantz and Parks (2002, Section 2.6).

We make three assumptions; the first two have already been motivated in Section 2.

Assumption 3.1. *The state evolution is*

$$\dot{x}(t)=\sum_{i=1}^{n}a_i(t)+h(x(t)),\qquad x(0)=x_0,$$

with $h(x)$ a real analytic function.

Assumption 3.2. *Utilities $u_i(x,a_i)$ are real analytic and satisfy $\frac{\partial u_i}{\partial x}\leq 0$ and $\frac{\partial^2 u_i}{\partial a_i^2}<0$.*

We add the following assumption on the boundary payoffs.

Assumption 3.3. *The boundary payoffs satisfy the inequalities $\beta_i(x_{\min})\geq u_i(x_{\min},a_i)/\rho$ and $\beta_i(x_{\max})\leq u_i(x_{\max},a_i)/\rho$ for all $a_i\in\mathcal{A}_i$.*

The third assumption, together with the fact that the $u_i(x,a_i)$ are decreasing in x, is used to show that the value functions $V_i(x)$ are decreasing in x: the state is a public bad.

if $x>\bar{x}$. But this value cannot be realised by a Markovian best response ϕ_1 of player 1, for if ϕ_1 were such that $V_1^{(\phi_1,\phi_2,\phi_3)}=V_1$, then $\phi_1=\phi$, and in particular $\phi_{1,-}(\bar{x})=1$, $\phi_{1,+}(\bar{x})=-1/2$, and $\mu=1/3$ at the discontinuity $\bar{x}$. However, the optimal trajectory-action pair for player 1 involves setting $a_1(t)=1$ and $\mu_1(t)=0$ for all t such that $x(t)=\bar{x}$, implying $V_1^{(\phi_1,\phi_2,\phi_3)}(\bar{x})<V_1(\bar{x})$, a contradiction. Loosely speaking, at $\bar{x}$ a jump point and a bang-bang point coincide. In terms of the concept introduced at the beginning of this section, the game is not 'strategically complete'. This example violates our assumption of $u(x,a_i)$ being monotonic in x, as well as Assumption 3.2 below.

The boundary payoffs only matter in situations where a player can take the state out of the state space, and so end the game; see Theorem 3.2 below. Often models can be set up in such a way that exit is not possible, as in the application in Section 4.

Theorem 3.1. *Given a collection $\mathcal{J}$ of jump points, for every i there is a shy and nowhere dense set $\mathscr{E}_{\mathcal{J},-i} \subset \mathscr{S}_{\mathcal{J},-i}$ such that the Markovian best response mapping $\mathscr{B}_i : \mathscr{S}_{\mathcal{J},-i} \backslash \mathscr{E}_{\mathcal{J},-i} \to \mathscr{S}_i$ is well-defined: for every strategy profile $\phi_{-i} \in \mathscr{S}_{\mathcal{J},-i} \backslash \mathscr{E}_{\mathcal{J},-i}$, player i has exactly one Markovian best response $\phi_i \in \mathscr{S}_i$.*

The proof, detailed in Online Appendices B–E, is structured as follows. First, direct arguments are used to show a number of properties of the value function V_i of player i: the value is decreasing and continuous everywhere except for a finite number of points. Then V_i is shown to be the unique viscosity solution to the Hamilton–Jacobi–Bellman equation of player i within the class of functions that possess these properties. To reach this result, we extend results of Barles et al. (2013, 2014) on optimal control problems with discontinuous dynamics to the situation that the dynamics are not controllable at a jump point.

Having characterised the value function, we build a regularity theory for V_i. This is a new contribution. We begin by showing that V_i is differentiable on a dense set. The points of differentiability yield initial points of trajectories of the state–costate equations. As the state–costate system is piecewise real analytic, this implies that V_i is real analytic on the projection of these trajectories to the state space. The trajectories end on steady states of the state–costate equations, of which there are potentially infinitely many. The case that the steady states accumulate on a jump point of ϕ_{-i} is shown to be shy. Steady states which have an accumulation point which is not a jump point of ϕ_{-i} have to be on the centre manifold of the accumulation point. A result of Aulbach (1986) implies that all points on this centre manifold are necessarily steady states, and that the centre manifold is real analytic. The value function is also real analytic on the projection of the centre manifolds on the state space. Apart from the shy cases, we end up with finitely many intervals of steady states and finitely many non-constant state trajectories, along which the value function is real analytic; consequently the best response belongs to $\mathcal{S}_i$. Uniqueness of the best response is clear from the construction, and the result is proved.

The theorem shows that our specifications of differential game and Markovian strategy space are fit for purpose: they ensure almost full strategic completeness, in the sense that each player will have a unique best response, in the same strategy space, to any profile of the other players' strategies in the complement of the exceptional shy set $\mathscr{E}_{\mathcal{J},-i}$, and that the best response generates optimal action schedules for all initial states.

Explicit sufficient conditions can be formulated for a strategy profile ϕ_{-i} to be not in the shy set $\mathscr{E}_{\mathcal{J},-i}$: these are of evident importance for applying Theorem 3.1. One such condition is given in Online Appendix E in Lemma E.10.

As the best response is piecewise real analytic, it can be characterised by classical conditions in the regions of analyticity, and by compatibility conditions at jump points. The second main result of this article, Theorem 3.2, formulates these. Its proof is given in Online Appendix F.

To state the theorem, we introduce the set

$$\mathcal{A}_{i,0} = \left\{ a_i \in \mathcal{A}_i : a_i + \mu g_{i,-}(x) + (1-\mu) g_{i,+}(x) = 0 \text{ for some } 0 \le \mu \le 1 \right\} \tag{3.15}$$

of *stabilising actions*: if we have $a_i = \phi_{i,-}(x) = \phi_{i,+}(x) \in \mathcal{A}_{i,0}$, then $0 \in F^\phi(x)$. At a boundary of an interval (x_1, x_2), where $x_1 < x_2$, we introduce the outward pointing normal $n(x)$ by setting $n(x_1) = -1$ and $n(x_2) = 1$. The set-valued function $F(x)$ *points into* (x_1, x_2) at a boundary point $x \in \{x_1, x_2\}$, if $n(x)\dot{x} \leq 0$ for all $\dot{x} \in F(x)$. If $x_1 = x_2$, the condition is void.

Theorem 3.2. *Let $\mathfrak{I}$ be a collection of jump points, and assume $\phi_{-i} \in \mathscr{S}_{\mathfrak{I},-i} \backslash \mathscr{E}_{\mathfrak{I},-i}$. Then $\phi_i = \mathscr{B}_i(\phi_{-i})$ if and only if the following hold.*

(i) Maximum principle: If $x \in \mathfrak{X} \backslash \mathfrak{I}$ and V_i^ϕ is differentiable at x, then

$$\phi_i(x) = a_i^*(x, (V_i^\phi)'(x)). \tag{3.16}$$

(ii) Monotonicity: V_i^ϕ is decreasing.

(iii) Value discontinuities: If V_i^ϕ is not continuous at x, then $x \in \mathfrak{I}$, $f_-^\phi(x)$ points into $(x_{\min}, x)$ and $F_+(x)$ into $(x, x_{\max})$.

(iv) Value at jump points: If $x \in \mathfrak{I}$, then $V_i^\phi(x) \geq \max_{a_i \in \mathcal{A}_{i,0}} u_i(x, a_i)/\rho_i$.

(v) Regularity at push–push steady states: If $x \in \mathfrak{I}$ is such that $f_-^\phi(x) > 0 > f_+^\phi(x)$, then $V_i^\phi(x)$ is differentiable at x.

The conditions can be interpreted. The Hamilton–Jacobi–Bellman equation of player i's value function reads as

$$\rho_i V_i(x) = u_i(x, a_i^*(x, V_i'(x))) + V_i'(x)(a_i^*(x, V_i'(x)) + g_i(x)), \tag{3.17}$$

while the value V_i^ϕ, obtained when the profile ϕ is played, satisfies

$$\rho V_i^\phi(x) = u_i^\phi(x) + (V_i^\phi)'(x)(\phi_i(x) + g_i(x)) \tag{3.18}$$

at points where ϕ is differentiable. If Condition (i) holds, then (3.18) implies that the value function V_i^ϕ satisfies the Hamilton–Jacobi–Bellman equation (3.17) at all points where V_i^ϕ is differentiable, that is, almost everywhere. This is the classical condition. By itself, it is insufficient to conclude that $V_i^\phi = V_i$ and that ϕ is the best response. In order to draw this conclusion, we have to show, first, that both V_i and V_i^ϕ are viscosity solutions of (3.17), and, second, that viscosity solutions of (3.17) are unique. The function of Conditions (ii)–(v) is to ensure that V_i^ϕ is a viscosity solution of (3.17).

Condition (ii) follows from the fact that the stock is a public bad, and says that there are no strategic incentives so perverse as to make the stock locally a good for player i.

The remaining conditions place restrictions on the best response at discontinuities of the other players' strategies. Condition (iii) says that a discontinuity in value is only possible at points where at least one of the other players' strategies is discontinuous, and which are such that player i is unable to control the dynamics back to the region of low stock if they ever end up on the high side of the discontinuity. Value discontinuities can occur at the boundary if it is impossible to exit the state interval. Condition (ii) implies that the value can only have downward (not upward) discontinuities.

Condition (iv) ensures that the value at a jump point is at least the value that can be obtained by stabilising the dynamics at that point.

Finally, Condition (v) follows from the fact that value is continuous at push-push states. If a player's best response is to be pushed strictly towards such a state, they end up at the same point whether approaching from the left or the right, and close to that point the continuity of the payoffs implies that the marginal value of the stock does not depend significantly on the direction of approach.

Our third result focuses on symmetric games, where all players have the same utility and the same discount rate, and on symmetric equilibria of these games. We leave the extension to non-symmetric equilibria to future work. Because of the symmetry, we drop the index i. The resulting ambiguity about the meaning of ϕ is resolved as follows: it will stand for the strategy of a generic player in ordinary expressions, and for the equilibrium profile when used as a superscript, as in f^ϕ.

The result gives a differential equation that has to be satisfied by equilibrium strategies in the interior $\mathring{\mathcal{A}}$ of the action space $\mathcal{A}$. We call $\phi(x)$ *interior* if $\phi(x)\in\mathring{\mathcal{A}}$. We introduce the *game Hamiltonian* in state–action variables

$$G(x,a)=u(x,a)-\frac{\partial u}{\partial a}(x,a)(Na+h(x)).$$

If ϕ is a Nash equilibrium strategy, then $G(x,\phi(x))$ gives the value of $H(x,(V^\phi)'(x))$, and hence of $\rho V^\phi(x)$, in the situation that $\Phi(x)=(N-1)\phi(x)$ and $\phi(x)$ is interior.

Theorem 3.3. *Let $\mathfrak{J}$ be a collection of jump points, and assume $\phi\in\mathscr{S}_\mathfrak{J}\setminus\mathscr{E}_\mathfrak{J}$. Then ϕ is a symmetric Nash equilibrium if and only if the following conditions hold.*

(i) Maximum principle: If $\phi(x)$ is interior, then

$$\frac{\mathrm{d}}{\mathrm{d}x}G(x,\phi(x))=-\rho\frac{\partial u}{\partial a}(x,\phi(x)). \tag{3.19}$$

Moreover for non-interior $\phi(x)$ we have that if $\phi(x)=a_{\min}$, then $(V^\phi)'(x)\leq -\frac{\partial u}{\partial a}(x,a_{\min})$, and if $\phi(x)=a_{\max}$, then $(V^\phi)'(x)\geq -\frac{\partial u}{\partial a}(x,a_{\max})$.

(ii) Monotonicity: V^ϕ is decreasing.

(iii) Value discontinuities: If V^ϕ is not continuous at x, then $x\in\mathfrak{J}$, $f_-^\phi(x)$ points into $(x_{\min},x)$ and $F_+(x)$ into $(x,x_{\max})$.

(iv) Value at jump points: If $x\in\mathfrak{J}$, then $V^\phi(x)\geq\max_{a\in\mathcal{A}_0}u(x,a)/\rho$.

(v) No push–push steady states: There is no point $x\in\mathfrak{X}$ such that $f_-^\phi(x)>0>f_+^\phi(x)$.

The differential equation (3.19) specifies a family of candidate Markov policies, to which the interior pieces of any piecewise-defined Markovian strategy must belong. The equation only depends on the fundamentals of the problem, and it can be solved without needing to know anything about the equilibrium.[27] This greatly simplifies the

27. A similar equation will not hold for non-interior points: knowledge of the value $V^\phi(x)$ is needed to determine the behaviour at jump points, and this depends on the entire equilibrium.

construction of any equilibrium.[28] Conditions (iii)–(v) then impose restrictions on how local pieces are connected together.

Note that even though our specification allows for situations in which utilities are computed as a convex combination of utilities of two different actions, when actions switch "infinitely often", Theorem 3.3 states that this outcome never arises in equilibrium. All symmetric MPEs (indeed, all MPEs) have well-defined payoffs also under classical solutions to the state dynamics.

We emphasise that Theorems 3.1 and 3.2 are of a new kind. All the conditions of Theorems 3.2 and 3.3 are moreover local, in terms of the state. They are used to construct the family of candidate Markov policies, using Equation (3.19), and to deduce global restrictions that rule out candidate strategies. We turn to an application to make this more concrete for the symmetric case.

4. APPLICATION

Our results can be used to find all symmetric Markov-perfect Nash equilibria, with finitely many jump points, of a differential game. We consider the canonical transboundary stock pollution game of van der Ploeg and de Zeeuw (1992). This model is characterised by a discount rate ρ and a quadratic flow utility

$$u(x,a) = \alpha a - \beta a^2/2 - \gamma x^2/2,$$

which are the same for all players. Here, a is the emission rate of a stock pollutant, and x is the pollution stock. Emissions generate benefits, the stock damages. The natural decay rate is proportional to the stock, $h(x) = -\delta x$, so that the dynamics of the pollution stock are

$$\dot{x}(t) = \sum_{j=1}^{N} a_j(t) - \delta x(t).$$

This linear–quadratic specification ensures the existence of a piecewise linear MPE that is defined on the whole state space.[29] Dockner and Ngo Van Long (1993) showed that the game also admits locally defined nonlinear equilibria of the type discussed, in a different context, by Tsutsui and Mino (1990).[30] However, as these locally-defined equilibria are not supported on the entire state space, the outcome is not defined in regions of the state space which could be reached with alternative strategies. It is then not clear in

28. Equation (3.19) is not new *per se*; similar equations have been used in applied work since at least Tsutsui and Mino (1990), and generalisations have been presented in various contexts (see e.g. Rincón-Zapatero et al., 1998). Note that in a discrete-time framework the Bellman equation cannot be solved without knowledge of the full value function, which depends on the equilibrium. In continuous time, the Hamilton–Jacobi–Bellman equation can be used locally to construct continuous candidate solutions. This benefit has been highlighted also in the recent literature on continuous-time macroeconomics (e.g. Achdou et al., 2014; Brunnermeier and Sannikov, 2016).

29. We often use "equilibrium strategy" and "MPE" synonymously.

30. Rubio and Casino (2002) show a further class of equilibria. Our methods subsume many of the "local MPEs" in the literature, showing how they are extended to globally defined discontinuous strategies.

TABLE 1
Calibration assumptions

Parameter	Quantity	Value	Unit
N	Number of players	5	
ρ	Discount rate	0.015	1/y
δ	Natural decay rate	0.001	1/y
x_0	Initial carbon stock	0.5	TtC
	Efficient steady state carbon stock	1.0	TtC
α/β	Business-as-usual emission rate	0.01	TtC/y
	Social cost of carbon today, first-best	400	T$/TtC

TABLE 2
Calibration results

Parameter	Value	Unit
α	678	T$/TtC
β	$339 \cdot 10^3$	T$ y/TtC^2
γ	1.953	T$/y TtC^2

TABLE 3
Natural units

Quantity	Natural unit	Expression
Time	417 y	$\sqrt{\beta/\gamma}$
Pollutant	0.833 TtC	$\alpha/\sqrt{\beta\gamma}$
Value	565 T$	$\alpha^2/\sqrt{\beta\gamma}$

what sense the purported equilibrium strategies can be said to be best responses (Rowat, 2007; Bernhard, 2024).

In illustrations, we use a stylised calibration in terms of atmospheric CO_2 pollution to show that this simple model produces quantitatively important results. We assume the players to be symmetric "major powers". Table 1 lists the calibration assumptions.

The unit abbreviations are 'y' for year, 'TtC' for 10^{12} metric tonnes of carbon, and 'T$' for 10^{12} US$. The low natural decay rate reflects the long persistence of CO_2 in the atmosphere; the efficient steady state refers to the steady state reached under full cooperation.[31] The business-as-usual emission rate is the optimal pollution level if pollution damages are not taken into account, that is, if $\gamma = 0$: given that the model is stationary, its value reflects emissions from fossil fuel use as of 2024. The calibrated values of the parameters are shown in Table 2.

Using these parameter values, we construct natural units of the model, given in Table 3. Expressing all quantities in multiples of these units leads to a model where $\alpha = \beta = \gamma = 1$, $\rho = 6.25$ and $\delta = 0.417$. As the natural time scale of the model is about four centuries, the natural discount rate ρ is not close to zero.

31. We calibrate current and efficient steady state carbon *stocks* to *cumulative emissions*. In reality, cumulative emissions are closely related to the degree of climate change, because of offsetting changes to the carbon-temperature relationship and the rate of decay of atmospheric carbon (Matthews et al., 2009). Our simple model should not be read too literally, but as an illustrative example.

We formulate our analytic results for N players, setting $\alpha=\beta=\gamma=1$, $\mathcal{A}=[0,1]$ and $\mathcal{X}=[0,N/\delta]$: the actions range between 0 and the bliss point, and the state space covers all states that are reachable from the initial state $x(0)=0$.

4.1. *Constructing the Markov-perfect Nash equilibria*

We use Theorem 3.3 to determine all piecewise continuously differentiable, symmetric Markov-perfect Nash equilibria, that is, symmetric policy rules $a=\phi(x)$ with finitely many discontinuities that give the emission flow as a function of the carbon stock. The starting point of the analysis is the differential equation (3.19). The pieces of piecewise strategies must be chosen from its solutions; thus our first task is to construct these.

In the transboundary pollution game, the Hamiltonian of an individual player is

$$H(x,p)=\max_a \left[a-a^2/2-x^2/2+p(a+\Phi(x)-\delta x)\right].$$

The optimal action satisfies

$$a^*(x,p)=\begin{cases}0 & \text{for } p\leq -1,\\ 1+p & \text{for } -1<p<0,\\ 1 & \text{for } p\geq 0.\end{cases}$$

For an interior action, using symmetry and $p=a^*(x,p)-1$, the game Hamiltonian is given by

$$G(x,a)=a-a^2/2-x^2/2+(a-1)(Na-\delta x).$$

We write Equation (3.19), which characterises equilibrium strategies away from jumps, as

$$\phi'(x)=\frac{(\rho+\delta)(\phi(x)-1)+x}{(2N-1)\phi(x)-(N-1)-\delta x}. \tag{4.20}$$

Any piecewise continuously differentiable symmetric Nash equilibrium ϕ has to satisfy the differential equation (4.20) locally, by Theorem 3.3(i). As the right hand side is real analytic if the denominator does not vanish, it follows that the solution ϕ is also real analytic. In particular, a piecewise continuously differentiable equilibrium is necessarily piecewise real analytic.

Figure 1 shows a number of local solutions $a=\phi(x)$ of Equation (4.20), drawn as solid curves; the figure focuses on the most interesting part of the state space. Candidate pieces of an equilibrium policy rule must coincide with such a solution. The dotted line indicates the points where the denominator of the right hand side of (4.20) vanishes. As almost all points on this *line of singularities* involve solutions with a vertical tangent, they cannot belong to any strategy in the strategy space $\mathscr{S}$, which requires at least one-sided differentiability.

The exception is the central singularity, where both numerator and denominator of (4.20) vanish: through this point two solutions pass that are both linear functions in a neighbourhood of the point. The graph of the decreasing solution, indicated by a

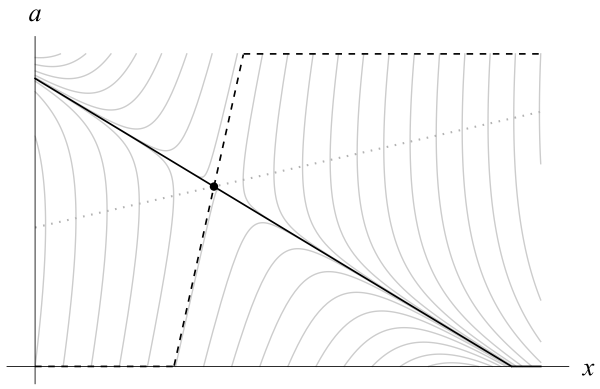

FIGURE 1
Local solutions of Equation (4.20). Indicated are the central singularity (black dot), the piecewise linear strategies $a = \phi_1(x)$ (solid) and $a = \phi_2(x)$ (dashed), as well as the line of singularities (dotted).

solid curve, intersects the lower bound $a = 0$: we extend this solution by setting its value equal to 0 for all states above the intersection point. We thus obtain a piecewise linear function ϕ_1, defined on the whole state space, which is the 'linear' Nash equilibrium strategy examined by van der Ploeg and de Zeeuw (1992).[32] We extend the increasing solution similarly to a piecewise linear function ϕ_2, assuming the values of the parameters ρ, δ and N to be such that the increasing solution intersects the lines $a = 0$ and $a = 1$ at states in the interior of the state space: this is the situation of our calibration.

With the family of solutions to Equation (3.19) at hand, we now start ruling out candidate equilibria by using conditions (ii)–(v) of Theorem 3.3. The task is to understand which policy rules ϕ, constructed from pieces of candidate solutions illustrated in Figure 1 and pieced together using the conditions, can be extended so that they cover the entire state domain.

Continuous equilibria. We first show that ϕ_1 is the unique continuous-strategy equilibrium. Assume that ϕ is a continuous Nash equilibrium strategy not equal to either ϕ_1 or ϕ_2. Then either $\phi(x) \geq \max\{\phi_1(x), \phi_2(x)\}$ everywhere, or $\phi(x) \leq \min\{\phi_1(x), \phi_2(x)\}$ everywhere: otherwise the graph of ϕ intersects the line of singularities not at the central singularity, and cannot be continued as the graph of a continuous function defined on the whole state space.

We infer that ϕ is either not emitting for small values of the state, or it is maximally emitting for large values of the state. That is, there is either an interval of positive length including $x = 0$ such that $\phi(x) = 0$ on that interval, or an interval of positive length including $x_{\max} = N/\delta$ such that $\phi(x) = 1$. We show that neither situation can obtain for an equilibrium strategy.

Consider first an equilibrium strategy ϕ that is not emitting for small values of the state. Clearly $V^{\phi}(0) = 0$, and, on an interval containing 0, Equation (3.18) reads as

$$\rho V^{\phi}(x) = -x^2/2 - \delta x (V^{\phi})'(x).$$

32. We remind the reader that we now consider only symmetric strategies ϕ, and the subscripts on ϕ_1, ϕ_2 indicate particular candidate strategies (and not player identity).

Solving this differential equation, we obtain $V^{\phi}(x) = -x^2/(4\delta+2\rho)$. However, Equation (3.16) requires that $(V^{\phi})'(x) \le -1$ close to $x=0$, which is not the case.[33]

The situation that ϕ is maximally emitting for large values of the state is even simpler. As V^{ϕ} is decreasing by Theorem 3.3(ii), we have $(V^{\phi})'(x) \le 0$ close to $x_{\max}$. But (3.16) implies that $(V^{\phi})'(x) \ge 0$ close to $x_{\max}$. Hence V^{ϕ} has to be constant there. Equation (3.18) however implies that $\rho V^{\phi}(x) = (1-x^2)/2$, which contradicts the constancy of V^{ϕ}. The argument shows more than claimed: it implies that an equilibrium strategy cannot be maximally emitting on any interval of positive length.

Hence there cannot be a continuous equilibrium strategy other than ϕ_1 or ϕ_2. Actually, the latter strategy is also disqualified, as it is both not emitting for small values and maximally emitting for large values. We therefore have obtained the following result.

Lemma 4.1. *The piecewise-linear strategy ϕ_1 is the unique continuous Nash equilibrium strategy of the transboundary pollution game.*

Discontinuous equilibria with continuous value. Having dealt with continuous equilibrium strategies, we use the family of interior candidate strategies, together with the local conditions imposed by Theorem 3.3, to construct global restrictions on discontinuous equilibrium strategies ϕ when the value V^{ϕ} is still continuous. To do this, we construct three regions and prove that no MPE strategy can pass through any of these. We then show that any point not in these regions can be part of an MPE strategy, and show how a globally-defined strategy can be constructed from such a point.

An easily verified observation is central to how candidate strategies can be pieced together:

Lemma 4.2. *The function $G(x,a)$ is convex in a, takes its minimum at the line of singularities, and satisfies $\rho V^{\phi}(x) = G(x,\phi(x))$ if $\phi(x)$ is interior and continuous at x.*

We begin with translating Condition (v) of Theorem 3.3, which excludes the possibility of push–push steady states (see Section 2.2) in equilibrium, into geometric terms.

Figure 2 is a modification of Figure 1. The dashed line indicates the steady state locus

$$Na - \delta x = 0.$$

At this locus total emissions are balanced by natural decay. Where the strategy involves emissions above this locus, the stock grows over time; where below, the stock diminishes. A push-push steady state is thus a situation connecting two candidate equilibrium strategy pieces such that the left piece is strictly above the locus, and the right piece is strictly below.

Let ψ_1 be the unique solution of Equation (4.20) whose graph is tangent to the steady state locus.[34] The shaded area indicates the union of the set of points under the steady state locus up to the tangency, and the points under the graph of ψ_1 (both strictly). This union will be called NE for "no emissions", as solutions of (4.20) in this region without push-push points necessarily do not emit for small values of the state. Formally:

33. In optimal control, such strategies would violate the transversality condition.
34. This solution was considered in some detail by Dockner and Ngo Van Long (1993).

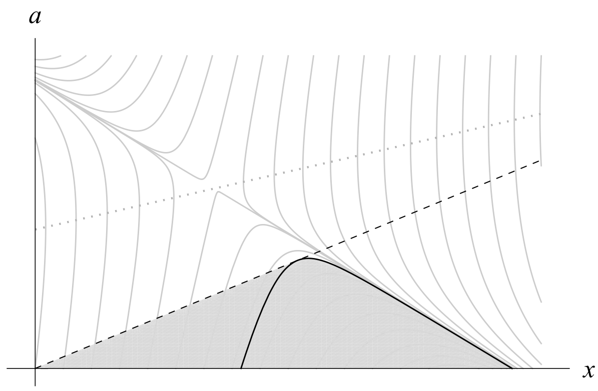

FIGURE 2
First forbidden region. Indicated are the region NE (gray), the steady state locus (dashed), the solution $a = \psi_1(x)$ tangent to it (solid), and the line of singularities (dotted).

Lemma 4.3. *The graph of a piecewise differentiable equilibrium strategy with continuous value function cannot intersect* NE.

Proof. Arguing by contradiction, we assume that ϕ is a piecewise differentiable equilibrium strategy with continuous value and that x_0 is a state such that $(x_0, \phi(x_0)) \in$ NE.

The region NE has the property that if $x_1 < x_0$ and ϕ is continuous in the interval $x_1 < x \le x_0$, then $(x_1, \phi_+(x_1))$ is also in NE.[35] If $\phi(x)$ is continuous for all states up to and including x_0, then necessarily $\phi(x)$ is not emitting at small values of the state. But this has already been shown to be impossible for an equilibrium strategy.

Hence ϕ has to have jump points below x_0. Let $\bar{x}$ indicate the largest of these. For all equilibrium strategies with continuous values, we have $\rho V^{\phi}(x) = G(x, \phi(x))$ if $\phi(x)$ is interior, see Lemma 4.2. The left and right hand limits of ϕ at $\bar{x}$ are moreover distinct at $\bar{x}$, that is $\phi_-(\bar{x}) \neq \phi_+(\bar{x})$, while by Condition (iii) of Theorem 3.3 their values are equal, so that

$$G(x, \phi_-(x)) = G(x, \phi_+(x)).$$

The game Hamiltonian $G(\bar{x}, a)$ is convex in a and takes its minimum at the line of singularities, which is above the steady state locus; the point $(\bar{x}, \phi_+(\bar{x}))$ is in NE, hence under the steady state locus and therefore under the line of singularities. Lemma 4.2 implies that $(\bar{x}, \phi_-(\bar{x}))$ is above the line of singularities, hence above the steady state locus. We conclude that $\bar{x}$ is a push-push steady state, which according to Theorem 3.3(v) yields a contradiction. ‖

The impossibility of push-push steady states has further implications: in Figure 3, we again indicate the steady state locus, but we added a second locus, the *conjugate* steady state locus, also indicated by a dashed line, of points $(x, a^{\dagger})$ which are such that if (x, a)

35. Recall that ϕ_+ refers to the right-hand limit.

is on the steady state locus, then $(x,a^\dagger)$ is above the line of singularities and

$$G(x,a^\dagger)=G(x,a).$$

Let ψ_2 be the unique solution of Equation (4.20) that is tangent to the conjugate steady state locus. The shaded region is the union of the set of points above the conjugate steady state locus from the tangency onwards, and the points above the graph of ψ_2. This region is called ME for "maximal emissions", as solutions of (4.20) in this region without push-push points necessarily emit maximally for large values of the state. Formally:

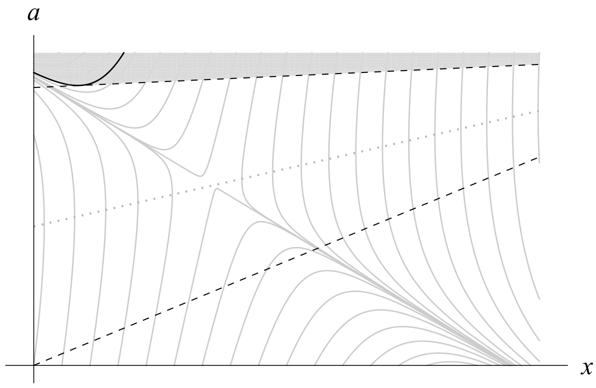

FIGURE 3

Second forbidden region. Indicated are the region ME (gray), the steady state locus and its conjugate (both dashed), the solution $a=\psi_2(x)$ tangent to the conjugate steady state locus (solid), and the line of singularities (dotted).

Lemma 4.4. *The graph of a piecewise differentiable equilibrium strategy with continuous value function cannot intersect* ME*.*

The proof is analogous to the proof of Lemma 4.3 and is omitted.

The third region that cannot be entered by equilibrium strategies is illustrated in Figure 4. It is

$$\mathrm{NC}=\left\{(x,a)\in\mathcal{X}\times\mathring{\mathcal{A}}:G(x,a)<G(x,\phi_1(x))\right\}.$$

The region is a subset of the interior, bounded by the graph of the unique continuous equilibrium strategy ϕ_1 and its conjugate $\phi_1^\dagger$, which is defined by the requirement that $G(x,\phi_1(x))=G(x,\phi_1^\dagger(x))$. We shall show that solutions of (4.20) in this region cannot be continued to globally defined continuous-value equilibrium strategies; hence NC is the "no continuation" region.

Lemma 4.5. *The graph of a piecewise differentiable equilibrium strategy with continuous value function cannot intersect* NC*.*

Proof. By contradiction: assume ϕ is a globally defined equilibrium strategy whose graph is in NC for some state x_0. Introducing $\Delta(x)=G(x,\phi(x))-G(x,\phi_1(x))$, this

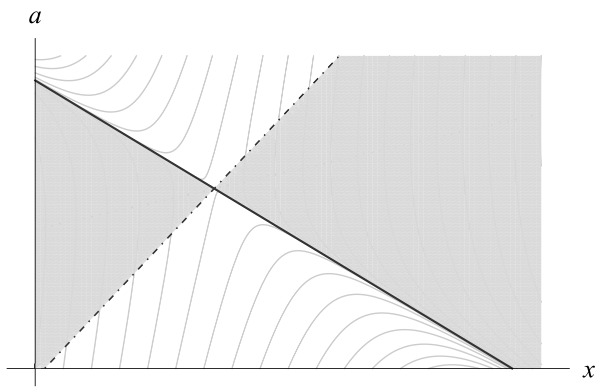

FIGURE 4
Third forbidden region. Indicated are the region NC (gray), the strategy $a=\phi_1(x)$ (solid), and its conjugate $a=\phi_1^\dagger(x)$ (dash-dotted).

is equivalent to assuming $\Delta(x_0)<0$. Let (x_s,a_s) denote the central singularity. As $a_s=\phi_1(x_s)$ minimises $G(x_s,a)$, it follows that $\Delta(x_s)\geq 0$.

We shall show that $\Delta(x)$ has to decrease as x moves from x_0 towards x_s. Equation (3.19) implies

$$\frac{\mathrm{d}}{\mathrm{d}x}\Delta(x)=\rho(\phi(x)-\phi_1(x)). \tag{4.21}$$

If ϕ jumps, the value of G, and hence of Δ, does not change: hence $\Delta(x)$ is continuous and piecewise differentiable in x, and Equation (4.21) holds everywhere except at finitely many points. It follows that Δ is decreasing to the left of the central singularity, where $\phi(x)<\phi_1(x)$, and increasing to the right of it. In particular $0\leq\Delta(x_s)<\Delta(x_0)<0$, which is a contradiction. ‖

Reformulating Lemma 4.5 gives one of our central results.[36]

Theorem 4.6. *The piecewise linear Markov-perfect equilibrium strategy ϕ_1 achieves the lowest possible payoff compared to all other continuous value equilibrium strategies.*

Thus the linear feedback MPE, a focus of the literature since Starr and Ho (1969), is Pareto-dominated by all other MPEs associated with a continuous value function. The graph of any locally defined candidate equilibrium ϕ yielding a lower value than the linear equilibrium is in the no-continuation region, and extending ϕ toward the central singularity of Equation (4.20) is not possible without losing the equilibrium property. At the central singularity, the linear strategy gives the lowest possible value of all candidate equilibrium pieces; no lower-valued equilibrium can extend across this point.

The *allowed* region $A=(\mathcal{X}\times\mathcal{A})\setminus(\mathrm{NE}\cup\mathrm{ME}\cup\mathrm{NC})$ (Figure 5) is the complement of the union of the forbidden regions. Our discussion up to this point can be summarised as

36. Compare this to Wirl (1996), who considers a situation where there are multiple continuous MPEs and who shows that the linear equilibrium achieves the lowest payoff of all continuous MPEs.

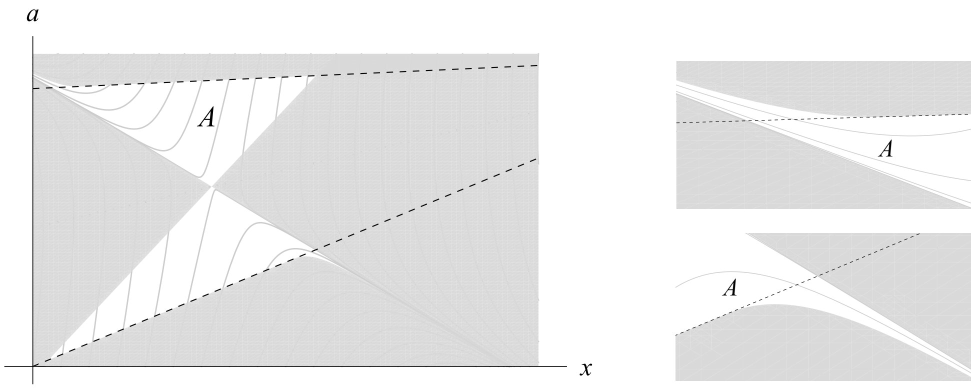

FIGURE 5

Left: the allowed region. Right: detailed view below the steady state locus (bottom) and above the conjugate steady state locus (top). Indicated are the region A (white), and the steady state locus and its conjugate (both dashed).

follows: the graph of a continuous-value equilibrium strategy is necessarily contained in the allowed region.

The next result characterises all equilibrium strategies with finitely many jump points.

Theorem 4.7. *A piecewise differentiable function ϕ with finitely many jump points is a continuous value Markov-perfect equilibrium strategy of the transboundary pollution game if and only if the following hold.*

(i) *Its graph is contained in the allowed region A.*
(ii) *At points where ϕ is differentiable and interior, Equation* (4.20) *holds.*
(iii) *The function $G(x,\phi(x))$ is continuous.*

The proof of this theorem is given in Appendix A.

Theorem 4.7 fully characterises all continuous-value MPEs with finitely many jump points. It has long been known that Equation (4.20) is necessary for an MPE. Our result adds the allowed region A to the necessary conditions, and shows that these three conditions are also sufficient.

The intuition underlying the welfare implications can be understood by introducing the effective discount factor for time t of a player as

$$D(t) \equiv \exp\left(-(\rho+\delta)t + \int_0^t (N-1)\phi'(x(s))ds\right).$$

This discount factor takes into account the effect of a marginal change in the carbon stock on other players future emissions. An individual player's costate equation

$$\dot{p}(t) = (\rho+\delta-(N-1)\phi'(x(t)))p(t) + x(t) = -\frac{\dot{D}(t)}{D(t)}p(t) + x(t),$$

can be written as

$$\frac{\mathrm{d}}{\mathrm{d}t}D(t)p(t) = D(t)x(t).$$

Integrating along the equilibrium path from $t=0$ to the next jump time $t=t_1$, or to $t_1=\infty$ if there is no jump, we obtain for the player's equilibrium cost of carbon $-p(t) = -V'(x(t))$ that

$$-V'(x_0) = \int_0^{t_1} D(t)x(t)dt + \lim_{t\uparrow t_1} D(t)(-V'(x(t))).$$

This is the present value of marginal damages up to time t_1, discounted by the effective discount factor $D(t)$, plus the discounted marginal damage value at t_1.

The linear equilibrium has no jumps, hence $t_1=\infty$. It is apparent from Figure 5 that the linear equilibrium strategy is the most accommodating strategy; i.e., $\phi'(x) \geq \phi_1'(x)$ for all ϕ. In other words, in the linear equilibrium, emissions are everywhere dynamic strategic substitutes: a marginal increase in carbon concentrations induces other players to lower their emissions, resulting in lower effective cost of carbon for the emitter. This is expressed through the lower effective discount factor. As all players reason in this way, collective pollution is high and eventually leads to high carbon stocks, harming welfare. The remaining non-linear equilibria involve less accommodation to carbon stocks than the linear equilibrium, so that their effective discount factor is always larger than its linear counterpart.

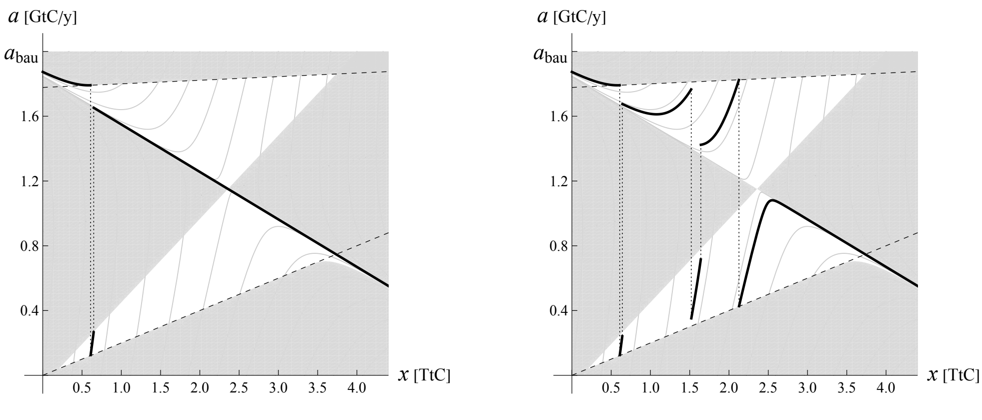

Figure 6

Two Markov-perfect Nash equilibria that are not piecewise linear. Indicated are the allowed region (white), an equilibrium strategy (solid, with dotted jumps), and the steady state locus and its conjugate (both dashed).

Figure 6 displays two MPEs, that is, two alternative strategies which are mutual best responses to themselves. The left panel displays a strategy which features high emissions for values of the carbon stock up to the critical level of 0.6 TtC. In this region, the equilibrium is Pareto-dominant, and carbon stocks grow rapidly. When the critical state is reached, emissions are reduced discontinuously and the state is stabilised at the critical level. If for some reason the stock grows above this steady state, the strategy prescribes all players to start increasing their emissions, and the economy starts drifting to higher carbon stocks. Upon reaching a second critical level of 0.65 TtC, emissions are increased discontinuously by all players, and the economy starts moving more rapidly towards a second steady state, with a carbon stock of 3.7 TtC. It is this threat of a substantial and rapid increase of the stock which ensures that the players are happy to stabilise at a lower level of the stock. The right panel displays a second equilibrium which behaves similarly at low carbon stocks, but in which there are two steady states at higher carbon stocks, and several discontinuities.

The low-carbon steady-state is thus semi-stable. To the right, $\phi'(x)$ is large and positive, so that the effective discount factor is large: if a player emits an additional unit, other players would increase their emissions rapidly, resulting in much higher costs for the original emitter: the equilibrium cost of carbon for the player is high. This reaction of the other players acts as a trigger response to deviations of the first player, which allows the equilibrium to stabilise such a small carbon stock with low emissions.[37]

Along the equilibrium trajectory approaching this steady state from below, the players understand that stabilisation at a small carbon stock is about to happen, and that the main effect of additional emissions is to bring forward the moment when emissions are discontinuously cut. The additional costs of the emission are to a large degree compensated by the fact that the low emission regime has been brought forward: hence the marginal cost of the additional amount of stock is low. After the emission cut, the beneficial effect of additional emissions is gone, and the trigger reaction described above of the other players sets in, all resulting in a discontinuous change of the player's cost of carbon.[38] Note that while low steady states can be supported by MPEs, these steady states are reached too fast, with suboptimally high emission rates, relative to the efficient outcome.

Equilibria with value discontinuities. Finally we turn to equilibrium strategies whose value function V^{ϕ} is not continuous. By Theorem 3.3(iii), a value discontinuity must be located at a point $\bar{x}$ at which the strategy is discontinuous, and the left hand limit $\phi_{-}(\bar{x})$ of the equilibrium strategy is below the steady state locus, while the right hand limit satisfies

$$a+(N-1)\phi_{+}(\bar{x})-\delta\bar{x}\geq 0$$

for all $a\in\mathcal{A}$, so that no player can individually take the state to the benign side of the discontinuity. This implies $\phi_{+}(\bar{x})\geq\delta\bar{x}/(N-1)$. If $\bar{x}$ is the smallest state at which there is a value discontinuity, then it is necessarily to the right of the unique tangency point of a solution of Equation (3.19) to the steady state locus. This gives a lower bound on the location of value discontinuities; in particular, any value discontinuity must be located in the right half of region NC, at a high carbon stock, above the steady-state of the linear equilibrium.

The right hand limit $\phi_{+}(\bar{x})$ is either above or below the line of singularities. If it is above that line, there has to be a second jump, this time with continuous value, to avoid the graph of ϕ entering the forbidden region ME of eventually maximally emitting strategies. To the right of this jump, the graph of ϕ is extended in the same manner as before. If

37. Sorger (1998) constructs a family of MPEs to a renewable resource game which almost stabilise the efficient steady state in a similar manner.

38. The stable equilibrium considered by Dockner and Ngo Van Long (1993) is also in the equilibrium set, extended to a global strategy. It does not have attractive welfare properties, as the steady-state carbon stock is quite high.

$\phi_+(\bar{x})$ is below the line of singularities, its graph can either be extended continuously for all larger states, or it may have continuous value jumps to points above the steady state locus, or discontinuous value jumps from points below this locus. See Figure 7.

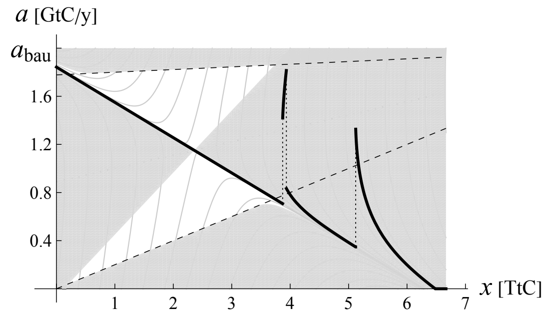

FIGURE 7

An equilibrium strategy with two value discontinuities. Legend as in Figure 6.

These equilibria are miscoordination traps, which sustain bad long-run outcomes, but from which no player can escape by unilateral action.

Equilibrium construction. Theorem 4.7 implies the following result.

Theorem 4.8. *Each point of the allowed region A is on the graph of an MPE strategy.*

Proof. The local solution of Equation (3.19) that passes through a given point can be continued for larger values of the state until either it intersects the conjugate steady state locus, from whence it jumps to the steady state locus, or it intersects the boundary of NC, from which it jumps to the linear strategy ϕ_1, or it intersects the line $a=0$, from which it is continued at that value. Continuing the solution for smaller values of the state is similar, only that then intersecting the steady state locus instead of its conjugate triggers a jump. Strategies constructed in this manner will have at most two jumps; however, more jumps can be inserted in the interior of the allowed region, respecting value matching and Lemma 4.2[39] ‖

Computation of particular equilibria is straightforward, requiring only the integration of Equation (4.20) between any points of discontinuity, and a continuous-value connection at these points.

Properties of Pareto-dominant equilibria. Theorem 4.8 allows us to construct the value envelope $\bar{V}$, which for every state gives the maximal value that can be supported by a Nash equilibrium. The solution to (3.19) tangent to the conjugate steady state locus is ψ_2; for states x below this tangency, we have that $\bar{V}(x)=V^{\psi_2}(x)$. Similarly, as the solution tangent to the steady state locus is ψ_1, for states x above that tangency we have $\bar{V}(x)=V^{\psi_1}(x)$. Between the two tangency points, $\bar{V}(x)$ is the value $G(x,a)$ for points (x,a) on the steady state locus. We recover the result of Schumacher et al. (2022).

39. Discontinuous actions do not need to satisfy a "smooth pasting" principle: for any player, the point of discontinuity is not endogenously chosen, but determined by the other players' strategies, taken as given.

Theorem 4.9. *There is no single Nash equilibrium strategy ϕ such that $\bar{V}(x) = V^{\phi}(x)$ for all x.*

Proof. The steady state locus between the two tangency states is not a local solution of (3.19). ‖

The welfare properties of MPEs are illustrated by comparing them with the business-as-usual solution, which we take as the emission path of a decision maker that completely ignores costs arising from pollution, and the fully cooperative solution. In Figure 8(a) we give the added value, relative to the business-as-usual solution, of the cooperative solution, the linear equilibrium, which is the lowest value a continuous-value MPE can achieve, and the value envelope, which gives for every state the highest MPE value that can be achieved when the dynamics commence at that state. Figure 8(b) gives the corresponding relative values, where business-as-usual corresponds to 0% and the cooperative solution to 100%. The linear equilibrium consistently achieves a bit more than half of the possible efficiency gain, whereas the value envelope reaches between 80% and 100% of the first best solution.

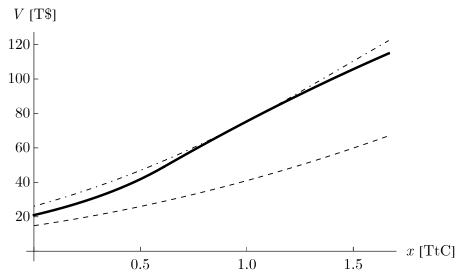

(a) Value differences

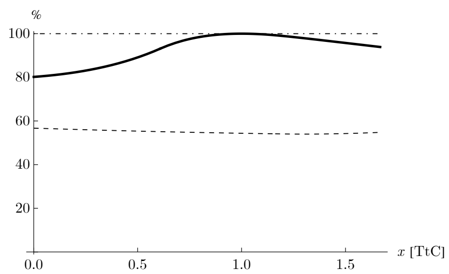

(b) Relative efficiency

FIGURE 8

Added value and relative efficiency, both with respect to business-as-usual, for different equilibrium strategies: cooperative (dash-dotted), best Nash (solid), linear Nash (dashed)

Stock evolutions and emissions are shown for two equilibria, the piecewise linear one and one that is Pareto optimal for the initial state $x_0 = 0.5$ TtC. The strategies are illustrated in Figure 9(a): the Pareto optimal strategy prescribes higher emissions until the pollution stock reaches 0.61 TtC, which happens in 13 years time, by which time emissions are cut to steady state level emissions of 0.12 GtC/y.[40]

Our results highlight that equilibrium multiplicity has important welfare implications. In particular, ensuring that a good equilibrium is played can lead to large welfare improvements. Climate negotiations can be seen as a forum for equilibrium choice (Dockner and Ngo Van Long, 1993). Coordination on a benign equilibrium may be more important than negotiating hard over who gets what fraction of the meagre surplus available to the players of a dynamic prisoner's dilemma.

40. Such a radical emission reduction could be challenging because of adjustment costs, reflecting unmodelled capital stocks. This makes an extension to more dimensions of particular interest.

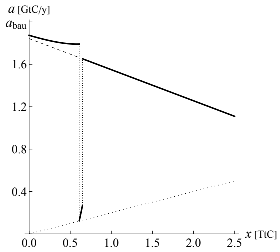

(a) Strategies

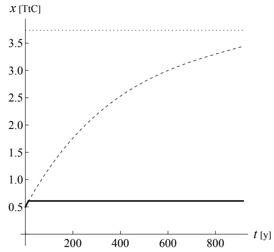

(b) Stock evolutions

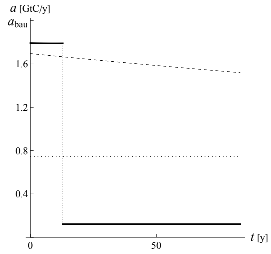

(c) Emissions

Figure 9

Strategies, stock evolutions and emissions for the equilibrium that is Pareto-dominant at initial state $x = 0.5$ TtC (solid), the linear equilibrium (dashed), and the steady state curve (dotted, panel 9(a)) and the steady state level of stock and emissions of the linear strategy (dotted, panels 9(b) and 9(c)).

5. CONCLUSION

We have shown how discontinuous strategies put differential games on a solid conceptual footing: best responses almost always exist and are unique. We have also given a necessary and sufficient condition for constructing best responses. Illustrating its use in a model of non-cooperative mitigation of climate change, we have constructed a large class of symmetric equilibria, all but one of them featuring discontinuous strategies. In other applications, with non-convex natural dynamics, well-defined continuous equilibria do not exist at all; our formulation allows the study of MPEs in such games also (Kossioris et al., 2008; Dockner and Wagener, 2014).

The importance of multiple equilibria in our application highlights the importance of equilibrium selection. Dockner and Ngo Van Long (1993) argued that climate negotiations could help states coordinate on a relatively benign equilibrium.[41] Our substantive message is twofold. First, in a non-cooperative world, gradual stabilisation of emissions gives poor strategic incentives and leads to bad outcomes; in our simple model, equilibria with gradual stabilisation have poor welfare properties. Second, more efficient long-run steady states are sustained in equilibrium by trigger-like (but Markovian) asymmetric responses, in which upward deviations from the benign steady state are punished by the players choosing collectively to slowly drift to a steady state with a higher carbon concentration. The efficient steady state is reached in finite time, with a discontinuous reduction in carbon emissions, suggesting that agreeing on radical climate policies might be better than agreeing on gradual reductions, specifically because radical policies provide better strategic incentives (similar to Dutta and Sundaram, 1993).[42] Such equilibria can be indexed by a “stabilisation target”, which triggers a “net zero” emission policy when reached. Thus, strategic leverage can be used to develop agreements with more effective incentives.

Some could find discontinuous strategies unrealistic. Indeed, while in the real world policy changes are almost always discontinuous, in many contexts large changes in policies are

41. As we have stressed, however, their treatment of multiple equilibria was problematic exactly for the reasons we address in the present paper. See also Dutta and Radner (2004).

42. van der Ploeg and Venables (2022) give further reasons for radical climate policies, based on strategic complementarities and multiple steady states.

rare. This could be seen to reflect adjustment costs. The appropriate response to such concerns would be to model these costs explicitly. However, the simplest model we can think of—two players emitting interacting flow pollutants—requires two state variables: the pollution flow for each player. This is already beyond our one-state framework, but we see no reason why such a set-up might not admit equilibria with coordinated jumps in the rates of change of the pollution flows. We must leave the extension of our results to multiple state variables for further work; this would of course greatly enlarge the range of potential applications.

Alternatively, could uncertainty, as in stochastic noise on the state evolution, or trembles on the actions, affect our results? We expect such noise to smooth out optimal responses near any discontinuities: the question is therefore whether every deterministic equilibrium is the vanishing noise limit of a family of stochastic equilibria. Of course, to answer this question fully necessitates a separate investigation. But we conjecture that the answer is affirmative, as our equilibria have been constructed on the basis of viscosity solutions, which were originally introduced to characterise the vanishing viscosity, that is, vanishing noise, limit of solutions to stochastic Hamilton–Jacobi–Bellman equations.

Finally, given that our methods have found several novel and important features of a well-studied model, we expect these methods to yield new insights in the numerous applications of differential games which exist in the literature, including games of incomplete information. Our results also allow, as well as motivate, the study of asymmetric equilibria.

A. PROOF OF THEOREM 4.7

Proof. The necessity of the conditions has been discussed in the run-up to the statement of Theorem 4.7. To show sufficiency, assume that ϕ satisfies the conditions.

Let x_1 be the state at which the graph of the linear equilibrium ϕ_1 intersects the steady state locus. By Condition (i), below the steady state locus and to the right of the line $x = x_1$, we have $\psi_1(x) \leq \phi(x) \leq \phi_1(x)$. Condition (iii) implies that ϕ is continuous, as there are no points in A above the steady state locus and to the right of the line $x = x_1$. In the interval $[a,b]$, where a and b are such that $\psi_1(a) = 0$ and $\phi_1(b) = 0$, we have $\phi'(x) < 0$ by Condition (ii). Hence there is a state $x_b \in [a,b]$ such that $\phi(x)$ is interior below x_b and $\phi(x) = 0$ above x_b.

Introduce the function

$$W(x) = G(x,\phi(x))/\rho = u(x,\phi(x))/\rho - \frac{\partial u}{\partial a}(x,\phi(x))(N\phi(x) - \delta x)/\rho. \tag{A.1}$$

We shall show that $W(x) = V^{\phi}(x)$ whenever $\phi(x)$ is interior. Equation (3.19) implies that

$$W'(x) = \frac{\mathrm{d}}{\mathrm{d}x} G(x,\phi(x))/\rho = -\frac{\partial u}{\partial a}(x,\phi(x)).$$

This allows us to rewrite Equation (A.1) as

$$\rho W(x) = u(x,\phi(x)) + W'(x)(N\phi(x) - \delta x).$$

This equation is the same as Equation (3.18), which is satisfied by the value function V^{ϕ}. Lemma A.1 implies that solutions to this equation are unique on all orbits of $\dot{x}(t) = N\phi(x(t)) - \delta x(t)$ that converge to a stable or a non-degenerate semi-stable steady state. As all trajectories of Equation (3.19) either intersect the steady state locus transversally or are non-degenerately, that is, quadratically, tangent to it, this is satisfied, and we conclude that $W(x) = V^{\phi}(x)$ and $\rho V^{\phi}(x) = G(x,\phi(x))$ for $x \leq x_b$, hence Condition (i) of Theorem 3.3 holds for interior $\phi(x)$.

We show Condition (i) also for non-interior $\phi(x)$. We have to show that for $x > x_b$ we have

$$(V^\phi)'(x) \leq -\frac{\partial u}{\partial a}(x,0) = -1.$$

The fact that x_b is the boundary point of the set of states where $\phi(x)$ is interior implies that $(V^\phi)'(x_b) = -1$ and $(V^\phi)''(x_b) \leq 0$. Arguing by contradiction, assume there is $x > x_b$ such that $(V^\phi)'(x) > -1$. Then there is $x_b < \xi < x$ such that $(V^\phi)'(\xi) = -1$ and $(V^\phi)''(\xi) \geq 0$.

On the interval $x > x_b$, Equation (3.18), which determines V^ϕ, takes the form

$$\rho V^\phi(x) = -\frac{1}{2}x^2 - \delta x (V^\phi)'(x).$$

Taking the derivative with respect to the state and rearranging yields

$$(V^\phi)''(x) = -(\rho+\delta)(V^\phi(x))'/(\delta x) - 1/\delta.$$

We find that

$$0 \leq (V^\phi)''(\xi) = (\rho+\delta)/(\delta\xi) - 1/\delta < (\rho+\delta)/(\delta x_b) - 1/\delta = (V^\phi)''(x_b) \leq 0,$$

which constitutes a contradiction. We conclude that Condition (i) holds.

The previous argument also shows that $V^\phi(x)$ is decreasing for non-interior $\phi(x)$. Using Equation (3.19), we find for interior $\phi(x)$ that

$$(V^\phi)'(x) = \frac{1}{\rho}\frac{\mathrm{d}}{\mathrm{d}x}G(x,\phi(x)) = \phi(x) - 1 < 0,$$

which shows that V^ϕ is decreasing everywhere: this shows Condition (ii) of Theorem 3.2.

Formally, we put $\beta(x_{\min}) = \infty$ and $\beta(x_{\max}) = -\infty$. The state space has been constructed such that the dynamics always point into it at the boundary, which takes care the value discontinuity at the boundary points: Condition (iii) holds there. Other value discontinuities are ruled out by the third assumption of the result to be proved, verifying Condition (iii) everywhere.

We verify Condition (iv) in the situation that $\phi_+(x) > \phi_-(x)$, the proof of the other situation being entirely similar. We write

$$\mathcal{A}_0 = \{a \in [0,1] : a - \delta x + (N-1)\phi = 0 \text{ for some } \phi \in [\phi_-(x), \phi_+(x)]\},$$

and we introduce $\bar{a}_- = \delta x - (N-1)\phi_-(x)$ and $\bar{a}_+ = \delta x - (N-1)\phi_+(x)$. Then

$$\max_{a\in\mathcal{A}_0} u(x,a) = \begin{cases} -\infty & \text{if } \bar{a}_- < 0 \text{ or } \bar{a}_+ > 1 \\ u(x,\bar{a}_-) & \text{if } 0 \leq \bar{a}_- \leq 1, \\ u(x,1) & \text{if } \bar{a}_+ \leq 1 < \bar{a}_-. \end{cases}$$

If $\bar{a}_- < 0$ or $\bar{a}_+ > 1$, there is nothing to prove.

Assume first that $0 \leq \bar{a}_- \leq 1$. If $\phi_-(x)$ is interior, we have

$$\begin{aligned} \rho V^\phi(x) - \max_{a\in\mathcal{A}_0} u(x,a) &= G(x,\phi_-(x)) - u(x,\bar{a}_-) \\ &= u(x,\phi_-(x)) - \frac{\partial u}{\partial a}(x,\phi_-(x))(N\phi_-(x) - \delta x) - u(x,\bar{a}_-). \end{aligned}$$

Noting that $N\phi_-(x) - \delta x = \phi_-(x) - \bar{a}_-$, and using the fact that $u(x,a)$ is concave in a, we obtain

$$\rho V^\phi(x) - \max_{a\in\mathcal{A}_0} u(x,a) = u(x,\phi_-(x)) + \frac{\partial u}{\partial a}(x,\phi_-(x))(\bar{a}_- - \phi_-(x)) - u(x,\bar{a}_-) \geq 0.$$

If $\phi_-(x)$ is not interior, it is zero and $\bar{a}_- = \delta x$. Above we have shown that then $(V^\phi)'(x) \leq -1 = \frac{\partial u}{\partial a}(x,0)$. Again using concavity of $u(x,a)$ in a, it follows that

$$\begin{aligned} \rho V^\phi(x) - \max_{a\in\mathcal{A}_0} u(x,a) &= u(x,0) + (V^\phi)'(x)(-\delta x) - u(x,\bar{a}_-) \\ &\geq u(x,0) + \frac{\partial u}{\partial a}(x,0)(\bar{a}_- - 0) - u(x,\bar{a}_-) \geq 0 \end{aligned}$$

We proceed by investigating the situation that $\bar{a}_- > 1$. If $\phi_-(x)$ is interior, we have

$$\rho V^\phi(x) - \max_{a\in\mathcal{A}_0} u(x,a) = \rho V^\phi(x) - u(x,1) = G(x,\phi(x)) - u(x,1)$$

$$\geq u(x,\phi_-(x)) - \frac{\partial u}{\partial a}(x,\phi_-(x))(N\phi_-(x)-\delta x) - u(x,\phi_-(x)) - \frac{\partial u}{\partial a}(x,\phi_-(x))(1-\phi_-(x))$$

$$= -\frac{\partial u}{\partial a}(x,\phi_-(x))((N-1)\phi_-(x)+1-\delta x) \geq 0.$$

At the first inequality, we have again used concavity; for the second inequality, we have used that $\frac{\partial u}{\partial a} \geq 0$ and $(N-1)\phi_-(x)+1-\delta x < 0$, which is a consequence of $\bar{a}_- > 1$.

In the non-interior case we have $\phi_1(x) = 0$ and $\bar{a}_- = \delta x > 1$. Then

$$\rho V^\phi(x) - \max_{a\in\mathcal{A}_0} u(x,a) = u(x,0) + (V^\phi)'(x)(-\delta x) - u(x,1)$$

$$\geq u(x,0) + \frac{\partial u}{\partial a}(x,0)(1-0) - u(x,1) \geq 0.$$

This concludes the verification of Condition (iv).

To verify Condition (v) of Theorem 3.3, the impossibility of push-push steady states, we remark that such a steady state would entail a jump from the part of the allowed region A above the steady state locus to the part below it. As the latter region is to the right of the unique tangency (x_1,a_1) of ψ_1 to the steady state locus, and the former to the left of the unique tangency (x_2,a_2) of ψ_2 to the conjugate steady state locus, it is sufficient to show that $x_1 > x_2$.

Let A and B be functions such that their graphs equal, respectively, the steady state locus and the conjugate steady state locus. That is, $A(x) = \delta x/N$ and $B(x) > A(x)$ satisfies $G(x,A(x)) = G(x,B(x))$ for all x.

We claim that if ϕ is a local solution of (3.19) that intersects the steady state locus at x_2, then $\phi'(x_2) > A'(x_2)$. Indeed, we have

$$\frac{\partial G}{\partial x}(x_2,A(x_2)) + \frac{\partial G}{\partial a}(x_2,A(x_2))A'(x_2) = \frac{\mathrm{d}}{\mathrm{d}x}G(x,A(x))\Big|_{x=x_2} = \frac{\mathrm{d}}{\mathrm{d}x}G(x,B(x))\Big|_{x=x_2}$$

$$= \frac{\partial G}{\partial x}(x_2,B(x_2)) + \frac{\partial G}{\partial a}(x_2,B(x_2))B'(x_2).$$

We have $\psi_2(x_2) = a_2 = B(x_2) > A(x_2) = \phi(x_2)$ and $\psi_2'(x_2) = B'(x_2)$. This implies

$$\frac{\partial G}{\partial x}(x_2,B(x_2)) + \frac{\partial G}{\partial a}(x_2,B(x_2))B'(x_2) = \frac{\partial G}{\partial x}(x_2,\psi(x_2)) + \frac{\partial G}{\partial a}(x_2,\psi(x_2))\psi'(x_2)$$

$$= \frac{\mathrm{d}}{\mathrm{d}x}G(x,\psi(x))\Big|_{x=x_2} = -\frac{\partial u}{\partial a}(x_2,\psi(x_2)) > -\frac{\partial u}{\partial a}(x_2,\phi(x_2)) = \frac{\mathrm{d}}{\mathrm{d}x}G(x,\phi(x))\Big|_{x=x_2}$$

$$= \frac{\partial G}{\partial x}(x_2,\phi(x_2)) + \frac{\partial G}{\partial a}(x_2,\phi(x_2))\phi'(x_2) = \frac{\partial G}{\partial x}(x_2,A(x_2)) + \frac{\partial G}{\partial a}(x_2,A(x_2))\phi'(x_2).$$

Here we have used Equation (3.19) in the third and fourth equality, and the inequality is a consequence of strict concavity of $u(x,a)$ with respect to a. We conclude that

$$\frac{\partial G}{\partial a}(x_2,A(x_2))(A'(x_2) - \phi'(x_2)) > 0,$$

and as $\frac{\partial G}{\partial a}(x,a) < 0$ along the steady state locus, the claim has been demonstrated.

Assume now that $x_1 \leq x_2$. As the graph of ψ_1 touches the steady state locus at (x_1,a_1) from below, it follows that all solutions intersecting the steady state locus at a point $(x,A(x))$ such that $x > x_1$ have a slope less than $A'(x)$. This contradicts the slope $\phi'(x_2)$ being larger than $A'(x_2)$. We conclude that $x_2 < x_1$, and that Condition (v) is satisfied. ‖

The following auxiliary result, used above, solves Hamilton–Jacobi–Bellman equations for the situation that utility and dynamics do not depend on an action variable. The proof is standard.

Lemma A.1. *Let g and v be real analytic, and let $\bar{x}$ be such that $g(\bar{x})=0$. Consider for $\rho>0$ the differential equation*

$$\rho W(x)-v(x)-W'(x)g(x)=0 \qquad a.e. \tag{A.2}$$

(i) *All solutions W of* (A.2) *are continuous close to $\bar{x}$ and satisfy $W(\bar{x})=v(\bar{x})/\rho$.*

(ii) *All solutions W are real analytic in $I\setminus\{\bar{x}\}$ for some open interval I containing $\bar{x}$, and $W'(x)$ tends to a finite or infinite limit as $x\to\bar{x}$.*

(iii) *If $g'(\bar{x})<0$, the equation has a unique and real analytic solution.*

(iv) *If $g'(\bar{x})=0$ and $g''(\bar{x})\neq 0$, the equation has a unique solution for $x<\bar{x}$ if $g''(\bar{x})>0$ and for $x>\bar{x}$ if $g''(\bar{x})<0$.*

Acknowledgements. For helpful comments, we thank Buz Brock, Guilherme Carmona, Larry Karp, Matt Mitchell, Rick van der Ploeg, Colin Rowat, Gerhard Sorger, Daniel Spiro, Georges Zaccour and Aart de Zeeuw, as well as seminar audiences at Aalto University, Colorado State University, GERAD, ifo Institute, the Swedish University of Agricultural Sciences, and the Universities of Amsterdam, Bologna, Grenoble, Oslo, Uppsala, and Utrecht. We also thank four anonymous referees as well as the editor for their constructive comments, which have helped greatly to improve the article.

Data availability statement

The code that has been used to generate the figures in this article is available on Zenodo at https://dx.doi.org/10.5281/zenodo.19660468.

REFERENCES

Achdou, Y., F. Buera, J.-M. Lasry, P.-L. Lions, and B. Moll (2014). Partial differential equation models in macroeconomics. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 372*(2028), 20130397.

Aulbach, B. (1986). Analytic center manifolds of dimension one. *Zeitschrift für Angewandte Mathematik und Mechanik 66*(3), 175–179.

Bardi, M. and I. Capuzzo-Dolcetta (2008). *Optimal control and viscosity solutions of Hamilton–Jacobi–Bellman equations.* Birkhäuser Boston.

Barles, G., A. Briani, and E. Chasseigne (2013). A Bellman approach for two-domains optimal control problems in $\mathbb{R}^N$. *ESAIM: Control, Optimisation and Calculus of Variations 19*(3), 710–739.

Barles, G., A. Briani, and E. Chasseigne (2014). A Bellman approach for regional optimal control problems in $\mathbb{R}^N$. *SIAM Journal on Control and Optimization 52*(3), 1712–1744.

Başar, T. and G. Olsder (1982). *Dynamic noncooperative game theory* (1st ed.). Philadelphia, PA: SIAM.

Bernhard, P. (2024). There is no known nonlinear markov perfect equilibrium strategies for the infinite horizon linear quadratic differential game. *Journal of Economic Theory 222*, 105927.

Board, S. and M. Meyer-ter Vehn (2013). Reputation for quality. *Econometrica 81*(6), 2381–2462.

Brunnermeier, M. and Y. Sannikov (2016). Macro, money, and finance: A continuous-time approach. In *Handbook of Macroeconomics*, Volume 2, pp. 1497–1545. Elsevier.

Dockner, E., S. Jørgensen, Ngo Van Long, and G. Sorger (2000). *Differential games in economics and management science.* Cambridge, UK: CUP.

Dockner, E. and Ngo Van Long (1993). International pollution control: Cooperative versus noncooperative strategies. *Journal of Environmental Economics and Management 25*(1), 13–29.

Dockner, E. and F. Wagener (2014). Markov perfect Nash equilibria in models with a single capital stock. *Economic Theory 56*(3), 585–625.

Dutta, P. and R. Sundaram (1993). The tragedy of the commons? *Economic Theory 3*(3), 413–426.

Dutta, P. K. and R. Radner (2004). Self-enforcing climate-change treaties. *Proceedings of the National Academy of Sciences 101*(14), 5174–5179.

Filippov, A. (1988). *Differential equations with discontinuous righthand sides.* Springer.

Fudenberg, D. and J. Tirole (1983). Capital as a commitment: Strategic investment to deter mobility. *Journal of Economic Theory 31*(2), 227–250.

Fudenberg, D. and J. Tirole (1991). *Game theory.* Cambridge, MA: MIT Press.

Hauser, D. N. (2024). Promoting a reputation for quality. *The RAND Journal of Economics 55*(1), 112–139.

Hunt, B., T. Sauer, and J. Yorke (1992). Prevalence: A translation-invariant "almost every" on infinite-dimensional spaces. *Bulletin of the American Mathematical Society 27*(2), 217–238.

Ivanov, R. (1990). Measurable strategies in differential games. *Mathematics of the USSR Sbornik 66*(1), 127–143.

Keller, G. and S. Rady (2015). Breakdowns. *Theoretical Economics 10*(1), 175–202.

Klein, N. and S. Rady (2011). Negatively correlated bandits. *The Review of Economic Studies 78*(2), 693–732.

Kossioris, G., M. Plexousakis, A. Xepapadeas, A. de Zeeuw, and K.-G. Mäler (2008). Feedback Nash equilibria for non-linear differential games in pollution control. *Journal of Economic Dynamics and Control 32*, 1312–1331.

Krantz, S. and H. Parks (2002). *A primer of real analytic functions* (2nd ed.). Boston: Birkhäuser.

Krasovskiĭ, N. and A. Subbotin (1988). *Game-theoretical control problems*. New York: Springer.

Matthews, H. D., N. P. Gillett, P. A. Stott, and K. Zickfeld (2009). The proportionality of global warming to cumulative carbon emissions. *Nature 459*(7248), 829–832.

Nash, J. (1951). Non-cooperative games. *Annals of Mathematics 54*(2), 286–295.

Rincón-Zapatero, J., J. Martínez, and G. Martín-Herrán (1998). New method to characterize subgame perfect Nash equilibria in differential games. *Journal of Optimization Theory and Applications 96*(2), 377–395.

Rowat, C. (2007). Non-linear strategies in a linear quadratic differential game. *Journal of Economic Dynamics and Control 31*(10), 3179–3202.

Rubio, S. and B. Casino (2002). A note on cooperative versus non-cooperative strategies in international pollution control. *Resource and Energy Economics 24*(3), 251–261.

Schumacher, J., P. Reddy, and J. Engwerda (2022). Jump equilibria in public-good differential games with a single state variable. *Dynamic Games and Applications 12*, 784–812.

Skiba, A. (1978). Optimal growth with a convex–concave production function. *Econometrica 46*, 527–539.

Sorger, G. (1998). Markov-perfect Nash equilibria in a class of resource games. *Economic Theory 11*(1), 79–100.

Starr, A. W. and Y.-C. Ho (1969). Nonzero-sum differential games. *Journal of Optimization Theory and Applications 3*(3), 184–206.

Sun, Y. (2024). A dynamic model of censorship. *Theoretical Economics 19*(1), 29–60.

Tsutsui, S. and K. Mino (1990). Nonlinear strategies in dynamic duopolistic competition with sticky prices. *Journal of Economic Theory 52*(1), 136–161.

van der Ploeg, F. and A. de Zeeuw (1992). International aspects of pollution control. *Environmental and Resource Economics 2*, 117–139.

van der Ploeg, F. and A. J. Venables (2022). Radical climate policies. World Bank Policy Research Working paper 10212.

Vinter, R. (2000). *Optimal control*. Boston: Birkhäuser.

Wagener, F. (2003). Skiba points and heteroclinic bifurcations, with applications to the shallow lake system. *Journal of Economic Dynamics and Control 27*, 1533–1561.

Wirl, F. (1996). Dynamic voluntary provision of public goods: Extension to nonlinear strategies. *European Journal of Political Economy 12*, 555–560.

Wirl, F. (2014). Uniqueness versus indeterminacy in the tragedy of the commons: A 'geometric' approach. In *Dynamic Optimization in Environmental Economics*, pp. 169–192. Springer.