Differential games of public investment

Niko Jaakkola* and Florian Wagener†

March 18, 2024

Abstract

We define a differential game of dynamic public investment with a discontinuous Markovian strategy space. The best response correspondence for the game is well-behaved: best responses exist and uniquely map almost all profiles of opponents' strategies back to the strategy space. Our chosen strategy space thus makes the differential game well-formed, resolving a long-standing open problem and allowing the analysis of a wider class of differential games and Markov-perfect equilibria. We provide a 'cookbook' necessary and sufficient condition for constructing the best response, and demonstrate its use with 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 Nash equilibria, and demonstrate that the best equilibria can yield a substantial welfare improvement over the equilibrium which previous literature has focused on. Our methods do not require specific functional forms.

*University of Bologna, Piazza Scaravilli 2, 40126 Bologna, Italy; CESifo. Email: niko.jaakkola@unibo.it

†Amsterdam School of Economics, University of Amsterdam, Roetersstraat 11, 1402 EJ Amsterdam, The Netherlands; Tinbergen Institute; CESifo. Email: f.o.o.wagener@uva.nl. Tel: +31 20 525 4034 (corresponding author)

The dynamic public goods game is an important economic problem which shows up in different settings, including joint investment projects between firms, allocation of effort among members of a team, harvesting renewable resources under common access, and non-cooperative mitigation of climate change. As with infinite-horizon dynamic games in general, these games typically admit the possibility of multiple equilibria—even under Markovian strategies, which condition each player's investment flow on the current state of the accumulating capital stock only. A natural question, with both positive and normative implications, is to ask what the entire set of Markov-perfect Nash equilibria (MPE) is. Except for special cases, the extant literature has not been able to address this issue.

We develop methods which enable progress on this question. We leverage the tractability of the continuous-time framework, which allows us to focus on the local properties of value functions without having to know the global properties.[1] Hence, we study a differential game,[2] making three contributions. First, we extend the space of Markovian strategies to include strategies that are discontinuous in the state variable and show that this yields a well-defined best response correspondence, putting a wide class of differential games on a conceptually sound foundation. Second, we give a concrete necessary and sufficient condition which can be used to construct a best response. We provide an example by solving for the set of symmetric MPE in a canonical model of non-cooperative mitigation of climate change. We obtain new results which are both intuitive and economically significant. While the application is primarily intended as an illustration of how our results can be used in applications, we believe the analysis is of interest in its own right.

We restrict the players to use Markovian strategies, or policy rules, so that the control schedule $a(t)$ is given as a function ϕ of the scalar state variable $y(t)$, that is $a(t) = \phi(y(t))$. The appropriate choice of strategy space—the set from which ϕ can be chosen—has long been an open question in the literature on differential games (Başar and Olsder, 1982; Fudenberg and Tirole, 1991; Dockner et al., 2000). To see why, note that computing the payoffs in a differential game requires the determination of a trajectory solving the state evolution equation $\dot{y}(t) = f(y(t), a(t))$. When using Markovian strategies, if the

[1]These benefits mirror recent advances in continuous-time macroeconomics (Achdou et al., 2014, 2022; Brunnermeier and Sannikov, 2016). There is also a large literature on repeated and dynamic games in continuous time: *inter alia*, on strategic experimentation (e.g., Keller et al., 2005; Klein and Rady, 2011), reputation in repeated games or dynamic games with imperfect monitoring (e.g. Sannikov, 2007; Faingold and Sannikov, 2011; Board and Meyer-ter Vehn, 2013), and on dynamic principal-agent problems (e.g., Sannikov, 2008; Cisternas, 2018).

[2]Dockner et al. (2000) and Başar and Zaccour (2018) contain extensive overviews of differential games.

function $f(x, q)$ and the strategies $\phi(x)$ are Lipschitz-continuous, then a unique classical solution trajectory exists. However, for a large class of models the optimal response to Lipschitz-continuous dynamics feature indifference initial states, often called Skiba points, at which there are multiple optimal solutions. In such a situation, the natural best response cannot be described as a Lipschitz-continuous function of the state (Skiba, 1978; Wagener, 2003), so that the natural best-response correspondence does not map the space of Lipschitz-continuous functions back to itself.

We therefore allow the players to use discontinuous Markovian strategies: player i's policy rule ϕ_i is selected from a space $\mathscr{S}_i$ of functions with a finite number of discontinuities. We then cannot apply the Picard-Lindelöf theorem on existence and uniqueness of solutions to differential equations. Hence, following Barles et al. (2013, 2014), we use a generalised solution concept to discontinuous dynamics, adapting the payoffs accordingly.

Our Theorem 1 shows that the resulting differential game is well-formed: the best response of player i to any profile $\phi_{-i} \in \mathscr{S}_{-i} \backslash \mathscr{E}$ of other players' strategies exists and can be described as a Markovian policy rule $\phi_i \in \mathscr{S}_i$. The exceptional set $\mathscr{E}$, for which the best response does not map to $\mathscr{S}_i$, is small: loosely speaking, it is shown to be a 'shy' set, that is, an infinite-dimensional analogue of a zero-measure set. We give a sufficient condition for identifying profiles of the other players' strategies that belong to $\mathscr{S}_{-i} \backslash \mathscr{E}$.

Our specification of the game is thus well-behaved, in that all strategy profiles induce a vector of payoffs; modulo exceptional cases, each player has a best response in $\mathscr{S}_i$; and each player can choose any strategy in $\mathscr{S}_i$ independently of the strategies chosen simultaneously by the other players, so that strategy profiles are a product set of individual strategies.[3] Our approach thus solves an issue which the literature on differential games has struggled with (and often ignored) and allows MPE in differential games to be interpreted in a way which is standard in game theory.

Theorem 2 is practical, giving necessary and sufficient conditions for a best response. While its proof is technical, the conditions are easy to apply. The result is general and does not require the use of particular functional forms.

[3]The existing literature often uses an admissibility criterion on strategy profiles, which implies that the set of strategies a player can choose from depends on strategies chosen *simultaneously* by other players, which sits awkwardly with the standard notion of simultaneous-move games (e.g., Dockner et al., 2000; Klein and Rady, 2011); although such an assumption may be more defensible when one of the players represents a fringe of non-strategic agents (Board and Meyer-ter Vehn, 2013).

In Section 6, we demonstrate how our results—in particular, Theorem 2—can be used to construct and analyse the fixed points of the best response correspondence, that is, the set of Nash equilibria in Markovian strategies. We construct the entire set of symmetric equilibria to a canonical model of non-cooperative mitigation of climate change (van der Ploeg and de Zeeuw, 1992; Dockner and Long, 1993). We adhere to the linear-quadratic framework used by these authors, even though our methods do not require it. We show that the linear equilibrium, by far the most commonly discussed equilibrium in the literature, is Pareto-dominated by all other symmetric equilibria with a value continuous in the state variable. This raises questions about the importance of the linear equilibrium from both positive and normative perspectives. We also characterise Pareto-dominant equilibria. These have a 'trigger'-like flavour, sustaining a favourable long-run level of the public good—a low atmospheric carbon stock—by means of asymmetric responses to deviations. A calibrated example shows that these equilibria can do much better than the linear equilibrium, or even sustain the socially optimal long-run carbon stock.

There is a large literature of other applications: we believe it is worthwhile to take a second look at these using our methods. Moreover, these methods can also shed light on asymmetric equilibria, something the existing literature has largely ignored.[4] An obvious question, left for future work, concerns extending our results to contexts with more than one state variable.

We use two primary tools in our analysis. The first is dynamic programming in the guise of the theory of viscosity solutions (Bardi and Capuzzo-Dolcetta, 2008). We apply viscosity theory to optimal control under discontinuous dynamics, building on the results by Barles et al. (2013, 2014), who consider exogenous discontinuities in dynamics. These methods allow us to construct the value function to a player's problem. We also rely on the theory of nonlinear dynamical systems to show that the best response is Markovian. Crucially, our ultimate goal is to understand equilibria, in which strategies—including any discontinuities—are endogenous. This means we cannot rule out complicated cases *a priori*.[5]

The present paper makes three broad contributions to the literature on dynamic games.

[4] In the context of strategic experimentation, Keller et al. (2005) and Keller and Rady (2015) develop asymmetric MPE. See also De Frutos and Martín-Herrán (2018).

[5] Viscosity solutions have been employed, in the context of games, by e.g. Sannikov (2008); Faingold and Sannikov (2011); Barilla and Gonçalves (2024).

First, we put the theory of Markov-perfect equilibria in a class of differential games—in particular, games in which the state variable can both decrease and increase—on a sound theoretical footing, as our specification makes the best-response correspondence well-behaved (at least when the state variable is a scalar). This issue has been an open problem for decades (Başar and Olsder, 1982; see also Fudenberg and Tirole 1991; Dockner et al. 2000). Our specification may also help in applications other than public investment.[6]

Second, our results demonstrate the power of continuous-time methods in deriving novel and general results in the analysis of dynamic investment games. Multiplicity of Markovian equilibria also obtains in discrete-time dynamic games. However, their general analysis is typically quite difficult. Our results flow from the fact that, in continuous time, the value function can be analysed and constructed using local information only.

Third, in terms of applications, our paper helps consolidate and clarify the literature on multiple MPE in differential games of public investment, starting with Tsutsui and Mino (1990) and Dockner and Long (1993), and continued by e.g. Dockner and Sorger (1996), Sorger (1998), Rubio and Casino (2002), Rowat (2007), and Dockner and Wagener (2014). Dutta and Sundaram (1993) construct, in a discrete-time context, an example with discontinuous Markovian strategies which sustains a long-run stock of a renewable resource above the socially optimal steady state. In our continuous-time framework, some equilibria share similar features.[7] However, our framework allows us to go further and to give a full and precise characterisation of the entire set of symmetric equilibria. We can also evaluate welfare outcomes, both in the steady state and in the transition to it—without needing to rely on "sufficiently low" values of the discount rate.[8] Our context is

[6]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 for the class of games we consider. We leave the question of whether our approach extends to alternative applications for future research.

[7]The equilibrium constructed by Dutta and Sundaram (1993) is not very robust: the high-stock steady state is unstable, in that a vanishingly small unexpected negative shock to the stock leads to the stock rapidly moving away to a low-stock steady state. In our application, the set of symmetric equilibria contains also such equilibria, but they can be seen as an atypical case; there exist more robust (semi-stable or stable) examples, supporting both more or less long-run exploitation of the public good (low degree of climate change) than what is socially optimal.

[8]Benhabib and Radner (1992) similarly consider a model of renewable resource exploitation; they show that a high public good stock can be enforced by trigger strategies based on past deviations (inferred implicitly by the state trajectory deviating from a prescribed one). We instead work with purely Markovian strategies. Dutta and Radner (2004) construct an equilibrium with purely Markovian

non-cooperative mitigation of climate change, a problem to which we give several novel results.

The paper proceeds as follows. Section 1 sets up the basic model. Section 2 previews the results of our application, to demonstrate the kinds of results our methods can obtain and to show that they matter. Section 3 gives an example of problematic situations our methods are designed to handle. Section 4 sets up the theoretical framework. The main results, existence and characterisation of the best response, are given in Section 5. We show how these are used in an application to climate change in Section 6. Section 7 wraps up. Proofs of the main results are technical and relegated to the Online Appendix.

1 Model

Time is continuous and runs to infinity: $t \in [0, \infty)$. The state space $\mathcal{X} = [x_{\min}, x_{\max}]$ is a compact interval of the set of real numbers $\mathbb{R}$. There are N players, indexed by $i \in \{1, \dots, N\}$.

Player i has access to an *action variable* $q_i \in \mathcal{Q}_i \subset \mathbb{R}$ through an *action schedule* $a_i : [0, \infty) \to \mathcal{Q}_i$. We assume the *control set* $\mathcal{Q}_i = [q_{i,\ell}, q_{i,u}]$ to be nonempty, convex and compact,[9] and action schedules $a_i(t)$ to be measurable functions. We collect actions and action schedules into vectors $q = (q_1, \dots, q_N)$ and $a(t) = (a_1(t), \dots, a_N(t))$. We write $q_{-i} = (q_1, \dots, q_{i-1}, q_{i+1}, \dots, q_N)$ and $q = (q_i, q_{-i})$. Similarly, we write $\mathcal{Q}_{-i} = \mathcal{Q}_1 \times \ldots \times \mathcal{Q}_{i-1} \times \mathcal{Q}_{i+1} \times \ldots \times \mathcal{Q}_N$. A *Markovian strategy* for player i is a map $\phi_i : \mathcal{X} \to Q_i$.

The state evolution depends on current state and actions, but not on calendar time: given a vector of action schedules a, the differential equation governing the state evolution is

$$\dot{y}(t) = f(y(t), a(t)). \tag{1}$$

A function $y : [0, \infty) \to \mathcal{X}$ satisfying $y(0) = x$ and (1) almost everywhere is a *classical trajectory*, and (y, a) a *classical trajectory–action pair*: these notions will be extended

enforcement of a desirable steady state, in a model with linear utility and persistent, irreversible climate mitigation efforts. Our model specification is very different; furthermore, we are able construct the entire set of equilibria. Nevertheless, this set contains some equilibria which are based on (very roughly) similar mechanisms.

[9] We could allow multivariate controls as in Dockner and Wagener (2014); our results would then require additional assumptions. Our key insights are best conveyed without such complications.

below. We distinguish between state and action variables x and q, and state trajectories y and action schedules a.

A function defined on an open set is *real analytic* if for any point in the set it can be represented, in a nonempty neighbourhood of the point, as a convergent power series with real coefficients. In this article we say that a function $\psi(x)$ is *piecewise real analytic*, if it is real analytic at all points, excepting a finite number of discontinuities, and such that the function and its derivative have finite limits as x approaches a discontinuity.

Assumption 1. *The function $f(x, q)$ is continuous, real analytic in x and q, and satisfies $f_{q_i} > 0$ everywhere.*

Thus, the action contributes to the growth of the state. The state variable is a public good, or public bad, in that the players' action variables reflect their contributions to investing in or disinvesting from it. In what follows, we focus on the latter case. The primitive of the payoffs is the flow felicity function:

Definition 1. *The* felicity *of player i when playing q_i at state x is $u_i(x, q_i)$.*

Assumption 2. *The felicity u_i is real analytic in x and q_i and satisfies $(u_i)_x < 0$ everywhere. Over time, felicity is exponentially discounted at a positive rate $\rho_i > 0$. For all x and q_{-i}, the set $\{(\eta_0, \eta_1) : \eta_0 \leq u_i(x, q_i),\ \eta_1 = f(x, q_i, q_{-i}),\ q_i \in \mathcal{Q}_i\}$ is convex.*

There is a unique maximiser $q_i = q_i^*(x, p, q_{-i})$ of $u_i(x, q_i) + pf(x, q_i, q_{-i})$ in Q_i. Moreover, there are real analytic functions $p_{i,\ell}(x, q_{-i})$ and $p_{i,u}(x, q_{-i})$, such that $q_i^* = q_{i,\ell}$ if $p \leq p_{i,\ell}$, $q_i^* = q_{i,u}$ if $p \geq p_{i,u}$, and q_i^* is a real analytic function of (x, p, q_{-i}) if $p_{i,\ell} < p < p_{i,u}$.

A unique maximiser clearly exists if, for example, one makes approriate assumptions on concavity.

Piecewise real analyticity covers the vast majority of parametrised models in the literature, usually specified using polynomial, rational, algebraic or elementary transcendental functions. In applications, functions f or u_i with singularities in their domain of definition are common; however, they are used for reasons of analytical convenience, as economic fundamentals rarely call for the presence of actual singularities. We do not need to resort to such specifications, and rule them out to sidestep unnecessary technical issues.

To be able to work with a compact state space, we have to specify the boundary behaviour. If $f(x_{\min}, a(t)) < 0$ or $f(x_{\max}, a(t)) > 0$, the state leaves the state space, the system is stopped, and player i receives a boundary value payoff $\beta_i(x)$.

Assumption 3. *The boundary payoffs satisfy* $\beta_i(x_{\min}) \geq \max_{q_i} u_i(x_{\min}, q_i)/\rho_i$ *as well as* $\beta_i(x_{\max}) \leq \min_{q_i} u_i(x_{\max}, q_i)/\rho_i$.

This assumption is used to derive that the state variable is a public bad: to see its necessity, note that if, for instance, the boundary payoff $\beta_i(x_{\max})$ is large, close to $x_{\max}$ the state might be a public good, as it allows the players to reach a high boundary payoff.

Let Θ denote the infimum of the set $\{t > 0 : y(t) \notin \mathcal{X}\}$ if that quantity is finite and ∞ otherwise, and introduce $\mathcal{T} = [0, \Theta]$. In the absence of discontinuities, the overall payoff is given by the sum of future discounted felicity, or

$$\int_0^{\Theta} \exp(-\rho_i t) u_i(y(t), a_i(t)) \, \mathrm{d}t + \exp(-\rho_i \Theta) \beta_i(y(\Theta)). \tag{2}$$

For the payoffs to be consistent with the fundamentals of the model when strategies can be discontinuous, we will require a richer description of the payoffs for situations in which there is no classical solution to the dynamics given by equation (1). We thus postpone the full payoff specification until Section 4.3.

The basic set-up is one of dynamic public investment, in terms of non-cooperative management of a stock pollutant. Adjusting the signs of the partial derivatives of the felicity function, or modifying the dynamics, will allow the model to be interpreted, for instance, as one of joint investment into a common project with depreciating capital, or as a model of renewable resource exploitation.

2 Motivation: non-cooperative climate policy

We first present a parametrised application, both to connect our work to previous literature, and to demonstrate concretely the kinds of results we can obtain using our methods.

Consider the model of the previous section, specifying $u_i(x, q_i) = \alpha q_i - \beta q_i^2/2 - \gamma x^2/2$ and $f(x, q) = \sum_{i=1}^{N} q_i - \delta x$. This is the canonical transboundary stock pollution model, often used for noncooperative climate change mitigation (van der Ploeg and de Zeeuw, 1992; Dockner and Long, 1993). The control takes values up to the bliss point: $q_i \in [0, \beta/\alpha]$.

We use the following stylised calibration to show that this simple model produces quantitatively important results. We assume the players are symmetric "major powers", choosing $N = 5$. The discount rate and the natural decay rate are respectively $\rho = 1.5\%$ and $\delta = 0.1\%$ per year. The initial state is 0.5 TtC, and the steady state of the socially optimal strategy is located at 1.0 TtC.[10] The 'business-as-usual' emission rate (not taking pollution damages into account) equals 0.01 TtC/y, reflecting emissions from fossil fuel use as of 2024.[11] Finally, the social cost of carbon in the linear equilibrium (see below), at the initial state, is 400T\$/TtC. This results in $\alpha = 181$ T\$/TtC, $\beta = 90.7 \cdot 10^3$T\$ y/TtC^2, and $\gamma = 0.522$ T\$/y TtC^2.

We are interested in Markovian strategies, which determine players' action schedules according to $a_i(t) = \phi_i(y(t))$ for policy rules ϕ_i. An MPE is a quintuple $\{\phi_i\}_{i\in\{1,\dots,5\}}$ such that each ϕ_i is a best response to ϕ_{-i}, starting from any initial state $x \in X$.[12]

It is well-known that, given the linear-quadratic specification, an MPE in symmetric piecewise linear strategies exists and can be computed using the 'guess-and-verify' method of equating coefficients (van der Ploeg and de Zeeuw, 1992). The single kink in this strategy occurs due to an active non-negativity constraint on the emission rate when the stock is high. This equilibrium strategy is shown as the dashed line in Figure 1(a).

The solid line in Figure 1(a) presents an alternative symmetric equilibrium strategy. This alternative strategy features higher emissions near the initial carbon stock $x = 0.5$ TtC. The implied trajectory features growth of the state until it reaches the level $x^* = 0.6$ TtC, after 13 years, at which point the per-player emissions discontinuously jump down, from 1.79 GtC/year to $\phi(q^*) = 0.12$ GtC/year (Figure 1(c)). At the latter level, aggregate emissions equal the decay of emissions, so that the trajectory is stabilised at x^*.

It is straightforward to compute the discounted value obtained by each player using the above strategies (Figure 2(a)). We take as our reference point a 'business-as-usual' (BAU)

[10] We calibrate current and long-term-optimal 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.

[11] Note that the 'business-as-usual' is not a Nash equilibrium.

[12] We will be more precise in defining the equilibrium in Section 4.

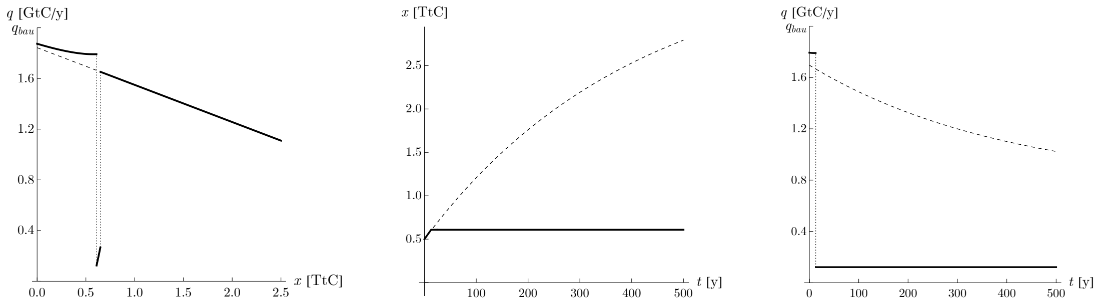

(a) Strategies (b) Stock evolutions (c) Emissions

Figure 1: Strategies, stock evolutions and emissions for the equilibrium that is Pareto-dominant at initial state $x = 0.5$ TtC (solid) and the linear equilibrium (dashed).

outcome, in which there is no policy to control carbon emissions. The total surplus value of implementing the social optimum, rather than the BAU, is 11.5 T\$. The surplus value of implementing the piecewise linear equilibrium is 6.4 T\$, achieving 56% of the socially optimal surplus. The value of implementing the discontinuous equilibrium strategy, on the other hand, is 9.9 T\$, or 87% of the available surplus. Thus, the non-linear strategy performs substantially better than the linear strategy and in fact gets fairly close to the social optimum. As this strategy is a symmetric MPE, the outcome is self-enforcing.

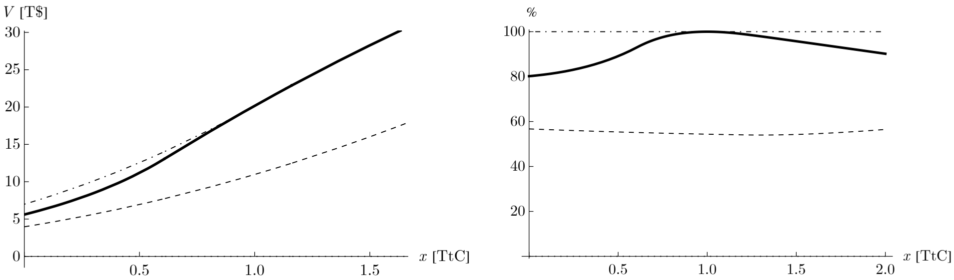

(a) Value differences (b) Relative efficiency

Figure 2: Added value and relative efficiency, both with respect to business-as-usual, for different equilibrium strategies: cooperative (dashed), best Nash (solid), linear Nash (dash-dotted)

In Section 6, we show that the discontinuous strategy of Figure 1(a) is in fact Pareto-dominant for the initial state $x = 0.5$ TtC, that is, it achieves the highest available payoff out of all symmetric MPE, of which there is a large set. Moreover, the piecewise linear equilibrium is the worst possible symmetric MPE for the same initial state. We can compute the highest possible symmetric MPE payoff for any other initial state: Figure 2(b) shows that these Pareto-dominant equilibria (all of which feature discontinuous strategies) achieve everywhere at least 80% of the socially optimal payoff surplus, achieve over 90% of the surplus over a wide region of the state space, and can also sustain the socially optimal steady-state stock at 1.0 TtC. This is a substantial improvement over the linear equilibrium, independent of the initial state. At $\rho = 1.5\%$ per year, these are not limiting results as the discount rate tends to zero.

The Pareto-dominant MPE sustains a relatively benign long-run steady-state carbon stock—indeed, the steady-state stock of 0.61 TtC is even below the socially optimal steady state (1 TtC), and well below the steady state of the piecewise linear equilibrium (3.74 TtC). This low long-run stock is sustained by a mechanism which resembles a trigger strategy. In a Markovian environment strategies cannot be conditioned on past actions, so the 'punishment' for emitting 'too much' is conditioned on the state becoming too high. Specifically, at the benign steady state, if the stock increases a little, the strategies specify that all players increase their emissions in response, thus pushing the stock further from the benign steady state. Moreover, there is a threshold at which emissions jump up discontinuously and thereafter follow the linear strategy, leading to rapid growth in the carbon stock and taking the economy to an adverse long-run steady state. It is this threat of a bad ultimate outcome which enforces the benign steady state.

We have above referred to Pareto-dominant symmetric equilibria. To make claims about Pareto-dominance, we must describe the entire set of symmetric equilibria. In Section 6, we show how our Theorem 2, stated in Section 5, can be used to obtain this set. A reader mainly interested in applying our results can skip Sections 3 and 4 and move to Section 6, consulting the main theorems in Section 5 when necessary.

To consider discontinuous strategies properly, we have to engage with technical issues related to the computation of state trajectories and payoffs. We illustrate these issues in the next section, and develop the tools required to tackle them in Section 4. These issues do not arise in the Pareto-dominant equilibrium shown above, which can be fully interpreted using the dynamics given by equation (1) and the payoff integral in equation (2). In general our Theorem 2 implies that the technical complications do not arise in equilibrium. As this is a result, not an assumption, our analysis must start with a more general set-up.

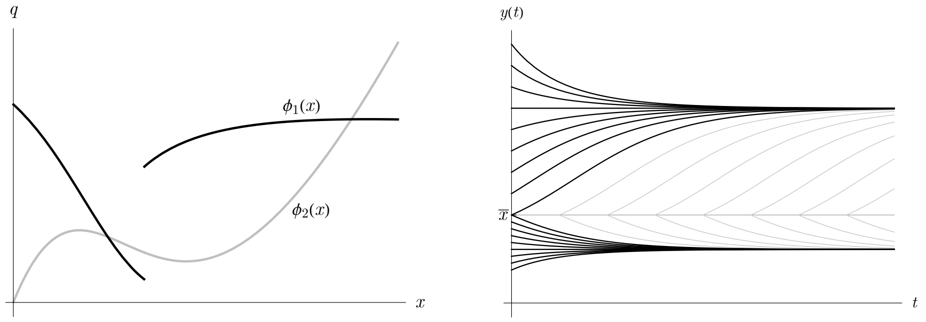

(a) Discontinuous best response ϕ_1 (b) Multiple trajectories starting at $\bar{x}$

Figure 3: Non-Lipschitz strategies generate a multiplicity of trajectories

3 Discontinuous strategies

In this section we demonstrate the technical issues with differential games that are solved by our approach.[13] Take the model of the previous section but with $N = 2$. Consider player 1 facing player 2's Markovian strategy ϕ_2. To compute the payoffs, given by equation (2), for any strategy ϕ_1, we need to determine the trajectory $y(t)$ induced by equation (1) and the pair (ϕ_1, ϕ_2); that is, the solution to

$$\dot{y}(t) = \phi_1(y(t)) + \phi_2(y(t)) - \delta y(t), \qquad y(0) = x \tag{3}$$

for any $x \in \mathcal{X}$. By the Picard-Lindelöf theorem, a unique classical solution is guaranteed to exist if the right-hand side $f^\phi(x) = f(x, \phi_1(x), \phi_2(x))$ is Lipschitz continuous in x. Given Assumption 1, this is guaranteed if ϕ_1 and ϕ_2 are both Lipschitz continuous.

Lipschitz continuity of player 1's best response ϕ_1 however fails to obtain in general. Take for instance $\phi_2(x) = bx - x^2/(1 + x^2)$, which is clearly Lipschitz continuous. Then player 1 faces an optimisation problem with concave-convex dynamics, the natural solution of which is known, for an open set of parameters, to feature an "indifference" or "Skiba" point $\bar{x}$ (Wagener, 2003). This is a discontinuity of the optimal policy function $q_1 = \phi_1(x)$ of player 1, see Figure 3(a). The resulting dynamics have two locally stable steady states x_1^s and x_2^s: which is optimally reached depends on the initial state. For $x < \bar{x}$, the optimal trajectory satisfies $y(t) \to x_1^s$, and for $x > \bar{x}$, $y(t) \to x_2^s$. The Markovian best response does not give rise to a unique solution trajectory at the initial point $x = \bar{x}$ (Figure 3(b)).

In other words, there exist perfectly standard optimisation problems such that the best

[13]See also Appendix B in Klein and Rady (2011).

Markovian response to Lipschitz continuous strategies may not exist in the space of Lipschitz continuous functions. In the following we allow the players to use discontinuous strategies. More precisely, the strategies can have a finite number of discontinuities ("jumps"). Everywhere else, they are real analytic, and the strategies and their first derivatives have finite one-sided limits when approaching a discontinuity. With the loss of Lipschitz continuity, both existence and uniqueness of solutions become problematic.

Two new situations arise. First, the 'push-push' situation, when $f^\phi_- \equiv \lim_{x\uparrow\bar{x}} f^\phi(x) > 0$, $f^\phi_+ \equiv \lim_{x\downarrow\bar{x}} f^\phi(x) < 0$ and $f^\phi(\bar{x}) \neq 0$. In the neighbourhood of $\bar{x}$, the natural solution is for the solution to reach $\bar{x}$ in finite time t_1 and remain there. However, there is no classical solution which satisfies the dynamics (3) for almost all $t > t_1$.

The second new situation is 'pull-pull', for which $f^\phi_- < 0$ and $f^\phi_+ > 0$. This arises in the situation of Figure 3(a). There are two classical solutions with initial state $\bar{x}$. In what follows, we will use a generalised solution concept to the dynamics, which also admits a continuum of 'irregular' natural solutions, indexed by a parameter $t_1 \geq 0$, such that $y(t) = \bar{x}$ for $0 \leq t \leq t_1$ and $y(t) \neq \bar{x}$ for $t > t_1$.[14] These are illustrated in Figure 3(b).

In what follows, we specify strategy spaces $\mathscr{S}_i$, extending the notion of a solution to the dynamics and adapting payoffs accordingly. We then show that best responses to opponents' strategies in $\mathscr{S}_{-i}$ almost always exist and belong to $\mathscr{S}_i$.

Our approach resolves a long-standing conceptual problem with differential games (discussed e.g. by Başar and Olsder, 1982). In the literature, the complications arising from the possibility of non-Lipschitz strategies are usually assumed away by either requiring strategies to be Lipschitz continuous, or with an admissibility requirement ruling out strategy *profiles* which lead to pull-pull or push-push dynamics. The former approach implies non-existence of a best response to many strategies available to the other players; the latter implies that the set of strategies player i can choose from depends on the strategies chosen by the other players. Our approach allows the space of admissible strategy profiles to simply be the product set of individual strategy spaces, as is standard in game theory, while allowing for non-Lipschitz strategies and ensuring that best responses exist.

[14]A continuum of classical solutions exists also if either of f^ϕ_- and f^ϕ_+ equals zero,

4 Markov-perfect Nash equilibrium

This section specifies the game. The specification is necessarily technical: we have to use a generalised solution concept for the dynamics, which necessitates a modification of the way payoffs are computed, as we cannot rely on uniqueness of solutions. Finally, we have to specify what "Markovian" strategies are in our context. In the specification we adopt, everything works in the expected (classical) fashion when the dynamics are continuous; it is the points of discontinuity which require special attention.

We start by describing the strategy spaces. Next, we set up an individual player's problem of optimal control of a differential inclusion. We consider the problem for general control and state trajectories. We then define what it means for a best response to be "Markovian" in our context. Finally, we present the equilibrium concept.

4.1 Markovian strategies

The players use Markovian strategies $\phi_i : \mathcal{X} \to \mathcal{Q}_i$, conditioning their actions on the current state variable only, in the following precise sense. To every strategy is associated an *adapted covering* $\mathscr{X}$ of $\mathcal{X}$ by a finite number of closed intervals that have non-empty and mutually disjunct interiors; restricted to one of these covering intervals, the function ϕ_i is real analytic on the interior of the interval. The restricted function and its derivative can be continuously extended to the covering interval. Informally, a function ϕ_i is constructed of sections of real analytic functions, but with the possibility of discontinuities where two adjacent sections are pieced together; at such an interface the value of ϕ_i is not defined, but one-sided limits and derivatives exist. The set of such strategies is denoted $\mathscr{S}_i$.

In the remainder of this subsection, we fix an $(N-1)$-tuple ϕ_{-i} of Markovian strategies. When using Markovian strategies, players set their actions as $a_i(t) = \phi_i(y(t))$. Given an action schedule a_i, and a strategy profile ϕ_{-i}, the system evolves according to

$$\dot{y}(t) = f_i(y(t), a_i(t)) = f(y(t), a_i(t), \phi_{-i}(y(t))), \tag{4}$$

where we introduced the dynamics $f_i(x, q_i) = f(x, q_i, \phi_{-i}(x))$ facing player i. In optimal control problems, it can occur that the optimal policy is described by a discontinuous Markovian policy rule. Hence, when there are several players present, the best response of a player may be a discontinuous Markovian strategy. A description of the game therefore has to take into account the possibility that discontinuous strategies are being played.

When Markovian strategies are not required to be continuous, f_i may have discontinuities, so that the evolution equation (4) may not have classical solutions, or may have a multiplicity of solutions. In the remainder of this section we generalise our notion of solution and describe how payoffs are adapted to that notion.

4.2 Coverings and dynamics

We partition the state space into J regions so that player i faces continuous dynamics within each region. The player will have a separate action schedule for each region; however, only action schedules in the neighbourhood of the current state—at most two—are "active" at a given moment. Definitions 4 and 5 are central: they define the dynamics at the boundary of two regions as a weighted average of the two active action schedules.

Some notation: the interior of a set S is denoted $\mathring{S}$; its closure $\bar{S}$; its boundary $\partial S = \bar{S} \backslash \mathring{S}$. Fix an integer $J > 0$ and points $x_{\min} = \bar{x}_0 < \bar{x}_1 < \bar{x}_2 < \ldots < \bar{x}_{J-1} < \bar{x}_J = x_{\max}$: the $\bar{x}_j$ are the possible locations of the discontinuities. Introduce a *covering* $\mathscr{X} = \mathscr{X}(\bar{x}_0, \bar{x}_1, \ldots, \bar{x}_J) = \{\mathcal{X}_j\}_{j=1}^J$ of $\mathcal{X}$, where the $\mathcal{X}_j$ are the closed intervals $[\bar{x}_{j-1}, \bar{x}_j]$. Then the state space is the union of the $\mathcal{X}_j$ as $\mathcal{X} = \bigcup_{j=1}^J \mathcal{X}_j$. For $j \in \{1, \ldots, J-1\}$, the *interface* $\mathcal{I}_j$ between $\mathcal{X}_j$ and $\mathcal{X}_{j+1}$ is the intersection $\mathcal{X}_j \cap \mathcal{X}_{j+1} = \{\bar{x}_j\}$. The set $\mathcal{I} = \bigcup_{j=1}^{J-1} \mathcal{I}_j$ is the union of all interfaces.

Let $\mathscr{F}^N_{\mathscr{X}}$ be the space of functions $\phi : \mathcal{X} \to \mathbb{R}^N$ with interface points given by $\mathscr{X}$. Precisely $\phi \in \mathscr{F}^N_{\mathscr{X}}$ if for each j, the restriction $\phi_{,j}$ of ϕ to $\mathring{\mathcal{X}}_j$ is a real analytic function and $\phi_{,j}$ and its derivative $\phi'_{,j}$ can be extended continuously to the closed interval $\mathcal{X}_j$. Two functions $\phi, \psi \in \mathscr{F}^N_{\mathscr{X}}$ are considered to be identical if they coincide on all open intervals $\mathring{\mathcal{X}}_j$. The space $\mathscr{F}^N$ is the union of all the $\mathscr{F}^N_{\mathscr{X}}$ for different number and locations of interfaces

$$\mathscr{F}^N = \bigcup_{J=1}^{\infty} \bigcup_{\bar{x}_0 < \ldots < \bar{x}_J} \mathscr{F}^N_{\mathscr{X}(\bar{x}_0, \ldots, \bar{x}_J)}.$$

In terms of these spaces, the set of full strategy profiles is given as

$$\mathscr{S} = \left\{ \phi \in \mathscr{F}^N \,:\, \phi(x) \in \mathcal{Q}_1 \times \ldots \times \mathcal{Q}_N \text{ for all } x \in \mathcal{X} \right\}.$$

The strategy spaces $\mathscr{S}_i$ for player i and $\mathscr{S}_{-i}$ for all players except player i are defined analogously, with $\mathcal{Q}_i$ respectively $\mathcal{Q}_{-i}$ replacing $\mathcal{Q}_1 \times \ldots \times \mathcal{Q}_N$. We also declare strategy spaces $\mathscr{S}_{\mathscr{X}} = \mathscr{S} \cap \mathscr{F}^N_{\mathscr{X}}$ with given interface points.[15]

Definition 2. *A* Markovian strategy *of player i is a function $\phi_i \in \mathscr{S}_i$. A* (full) strategy profile *is an N-tuple of Markovian strategies $\phi = (\phi_1, \ldots, \phi_N) \in \mathscr{S}$. The strategy profile of all players except player i is denoted $\phi_{-i} = (\phi_1, \ldots, \phi_{i-1}, \phi_{i+1}, \ldots, \phi_N) \in \mathscr{S}_{-i}$.*

The *local dynamics* for player i are given, for $j \in \{1, \ldots, J\}$ and $x \in \mathcal{X}_j$, by

$$f_{i,j}(x, q_{i,j}) = f(x, q_{i,j}, \phi_{-i,j}(x)),$$

where $q_{i,j}$ is the *local action.* Note that $f_{i,j}$ is conditional on $\phi_{-i,j}$; we do not explicitly indicate this in order to avoid notational clutter.

We introduce the sets $F_{i,j}(x) = f_{i,j}(x, \mathcal{Q}_i)$ for $x \in \mathcal{X}_j$ and $F_{i,j}(x) = \emptyset$ for $x \in \mathcal{X} \backslash \mathcal{X}_j$. Then we define the set-valued map $F_i : \mathcal{X} \rightsquigarrow \mathbb{R}$ in terms of the $F_{i,j}(x)$ as

$$F_i(x) = \overline{\text{co}} \left(\bigcup_{j=1}^{J} F_{i,j}(x) \right).$$

Using F_i we can formulate the announced extension of our original notion of trajectory; we consider trajectories to be solutions to a diffential inclusion.

Definition 3. *Given a strategy profile $\phi_{-i} \in \mathscr{S}_{-i}$, a* (state) trajectory *of F_i with initial state $x \in \mathcal{X}$ is an absolutely continuous function $y : \mathcal{T} \to \mathcal{X}$ such that $y(0) = x$ and $\dot{y}(t) \in F_i(y(t))$ for almost all $t \in \mathcal{T}$.*

Formulating the dynamics in terms of F_i hides the actions. We give an equivalent formulation that shows them. The *indicator function* $\mathbf{1}_S$ of a set S is given as $\mathbf{1}_S(x) = 1$ if $x \in S$ and $\mathbf{1}_S(x) = 0$ if $x \notin S$.

Definition 4. *The* dynamics *for player i are a function $\mathbf{f}_i : \mathcal{X} \times \mathcal{Q}_i^J \to \mathbb{R}$ given as*

$$\mathbf{f}_i(x, q_i) = \begin{cases} \sum_{j=1}^{J} \mathbf{1}_{\mathcal{X}_j}(x) f_{i,j}(x, q_{i,j}) & \text{if } x \in \mathcal{X} \backslash \mathcal{I}, \\ \mu_{i,j}(q_i) f_{i,j,-} + (1 - \mu_{i,j}(q_i)) f_{i,j,+} & \text{if } x = \bar{x}_j \in \mathcal{I}, \end{cases}$$

[15]For any current state $y(t)$, at most two components of ϕ are active: the single component $\phi_{,j}$ when $y(t) \in \mathring{\mathcal{X}}_j$, or components $\phi_{,j}$, $\phi_{,j+1}$ when $y(t) \in \mathcal{I}_j$.

where $(f_{i,j,-}, f_{i,j,+}) = (f_{i,j}(\bar{x}_j, q_{i,j}), f_{i,j+1}(\bar{x}_j, q_{i,j+1}))$ *and* $\mu_{i,j}(q_i) = f_{i,j,+}/(f_{i,j,+} - f_{i,j,-})$ *if* $f_{i,j,+}f_{i,j,-} \leq 0$ *and* $(f_{i,j,+}, f_{i,j,-}) \neq (0,0)$*, and* $\mu_{i,j}(q_i) = 0$ *otherwise.*

Note that $\mu_{i,j}(q_i)$ is chosen such that $\mathbf{f}_i(\bar{x}_j, q_i) = 0$ for the push-push situation $f_{i,j,-} \geq 0$ and $f_{i,j,+} \leq 0$ as well as for the pull-pull situation $f_{i,j,-} \leq 0$ and $f_{i,j,+} \geq 0$.

Definition 5. *An* action schedule *of player* i *is a vector-valued function*

$$a_i(t) = (a_{i,1}(t), \ldots, a_{i,J}(t)),$$

with $a_{i,j} \in L^\infty(0,\infty;\mathcal{Q}_i)$ *the* local action schedules*. The set of action schedules of player* i *is denoted* $\mathcal{A}_i$*. If* y *is a state trajectory such that*

$$\dot{y}(t) = \mathbf{f}_i(y(t), a_i(t)) \tag{5}$$

for almost all $t \in \mathcal{T}$*, then* (y, a_i) *is called a* trajectory–action pair*.*

The local action schedule $a_{i,j}$ is the active action schedule if the trajectory y is in the relevant part $\mathcal{X}_j$ of the state space.

The next proposition is a selection result stating that every state trajectory is generated by some action schedule. It is a direct corollary of Barles et al. (2013, Theorem 2.1).

Proposition 4.1. *If* y *is a state trajectory, there exists an action schedule* a_i *such that* (y, a_i) *is a trajectory–action pair.*

The following result is a converse to Proposition 4.1: every action schedule generates a trajectory. The argument is straightforward and therefore omitted.

Proposition 4.2. *For every* $\phi_{-i} \in \mathscr{S}_{-i}$*,* $x \in \mathcal{X}$*, and* $a_i \in \mathcal{A}_i$*, there is a state trajectory* y *with initial state* x *such that* (y, a_i) *is a trajectory–action pair.*

Specifying an action schedule and an initial state does not uniquely determine a state evolution: if the initial state is at an interface, and the two active actions are pulling the state away from the interface, then both a trajectory that remains at the interface for a positive amount of time and a trajectory that veers away immediately are valid state trajectories. Such a pull-pull situation that goes on for a positive amount of time is inherently unstable and would be immediately resolved by the slightest perturbation. We call trajectories that do not display this behaviour 'regular'.

Definition 6. *Let* (y, a_i) *be a trajectory–action pair. If for almost all* $t \in \mathcal{T}$ *such that* $y(t) \in \mathcal{I}_j$ *for some* $j \in \{1, \ldots, J-1\}$ *we have*

$$f_{i,j}(y(t), a_{i,j}(t)) \geq 0 \quad \textit{and} \quad f_{i,j+1}(y(t), a_{i,j+1}(t)) \leq 0, \tag{6}$$

then the trajectory–action pair is called regular*.*

Given x*,* a_i *and* ϕ_{-i}*, the set of all trajectories* y *such that* (y, a_i) *is a trajectory–action pair, respectively a regular trajectory–action pair, is denoted* $\mathcal{Y}_{x,a_i,\phi_{-i}}$*, respectively* $\mathcal{Y}^{\text{reg}}_{x,a_i,\phi_{-i}}$*.*

4.3 Payoffs and boundary conditions

We need to specify the flow payoff at a boundary of two regions. We take it to be a weighted average of the flow payoffs in the neighbourhood of the current state, using the weight associated with the dynamics in the previous section.

Specifically, we define for $j \in \{1, \ldots, J\}$ and $x \in \mathcal{X}_j$ the *local flow payoffs* $u_{i,j}(x, q_{i,j}) = u_i(x, q_{i,j})$, which are equal to the flow payoffs if $x \in \mathring{\mathcal{X}}_j$. At an interface, the payoffs are given by a weighted average of the left hand and the right hand payoff, where the weights are the same as for the dynamics (see Definition 4).

Definition 7. *The* flow payoff *for player* i *is a function* $\mathbf{u}_i : \mathcal{X} \times \mathcal{Q}_i^J \to \mathbb{R}$ *defined as*

$$\mathbf{u}_i(x, q_i) = \begin{cases} \sum_{j=1}^{J} \mathbf{1}_{\mathcal{X}_j}(x) u_{i,j}(x, q_{i,j}) & \textit{if } x \in \mathcal{X} \backslash \mathcal{I}, \\ \mu_{i,j}(q_i) u_{i,j}(x, q_{i,j}) + (1 - \mu_{i,j}(q_i)) u_{i,j+1}(x, q_{i,j+1}) & \textit{if } x = \bar{x}_j \in \mathcal{I}. \end{cases}$$

To gain some intuition for the payoff specification, note that in a push-push situation, one can think of the state spending 'more time' on the side of the discontinuity towards which it faces greater 'pressure' by the dynamics, and the flow utility on this side of the discontinuity thus receives greater weight. Such an outcome would be generated by a discrete-time approximation of the dynamics, with short time intervals.

Total welfare is integrated discounted felicity:

Definition 8. *Given a trajectory–action pair* (y, a_i) *with initial state* x*, the* total payoff *from the pair for player* i *is given by*

$$U_i(y, a_i) = \int_0^{\Theta} \exp(-\rho_i t)\mathbf{u}_i(y(t), a_i(t))\,\mathrm{d}t + \exp(-\rho_i \Theta)\beta_i(y(\Theta)).$$

The value *at the initial state* x *of the profile* ϕ_{-i} *to player* i *is*

$$V_i(x) = \sup_{\mathcal{A}_i} \sup_{\mathcal{Y}_{x,a_i,\phi_{-i}}} U_i(y, a_i)$$

where the first supremum is taken over the action schedules a_i*, and the second over the set of trajectories* y *for player* i*, given* x*,* a_i*, and* ϕ_{-i}*.*

The regular value V_i^{reg} *is defined analogously, with the set* $\mathcal{Y}_{x,a_i,\phi_{-i}}$ *of trajectories replaced by the set* $\mathcal{Y}^{\mathrm{reg}}_{x,a_i,\phi_{-i}}$ *of regular trajectories.*

An action schedule a_i *for which the supremum is realised is called a* best response of player i*. The set of all trajectories, respectively all regular trajectories, that are associated to a best response is denoted* $\mathcal{Y}^{*}_{x,\phi_{-i}}$*, respectively* $\mathcal{Y}^{\mathrm{reg},*}_{x,\phi_{-i}}$*.*

As a consequence of Proposition 4.2, the value V_i is finite for all x. An important technical result will be to show that the condition $(u_i)_x < 0$ implies that V_i and V_i^{reg} are identical.

4.4 Markovian best responses and MPE

We next define what it means for a best response to be "Markovian" in our setting. Informally, a best response of player i is Markovian if it can be described fully by a function ϕ_i of the current state only. Where strategies are discontinuous, the associated flow payoff is described by left and right limits of the Markovian strategy, again as a weighted average of the limiting flow payoffs, with a weight consistent with the dynamics. In other words, active action schedules are continuous even when arriving at a boundary between two regions. We need to show that best responses are Markovian, as this is not obvious (Barles et al., 2013).

Let therefore a full strategy profile ϕ be given, as well as a covering $\mathscr{X}$ adapted to it.

The evolution equation reads then as

$$\dot{y}(t) = f(y(t), \phi(y(t))) \equiv f^{\phi}(y(t)). \tag{7}$$

As before, solutions are not well-defined at discontinuities of f^{ϕ}. A solution, or state trajectory, of (7) is henceforth defined as an absolutely continuous function $y(t)$ that satisfies the differential inclusion

$$\dot{y}(t) \in F^{\phi}(y(t)) \tag{8}$$

almost everywhere, where the set-valued map F^{ϕ} is given as $F^{\phi}(x) = \{f^{\phi}(x)\}$ if x is not an interface point and $F^{\phi}(x) = \overline{\mathrm{co}}\{f^{\phi}_{,j}(x), f^{\phi}_{,j+1}(x)\}$ for $x = \bar{x}_j$.

The differential inclusion reduces to (7) if $y(t)$ is not at an interface point. At an interface $y(t) = \bar{x}_j$ we have

$$\dot{y}(t) = \mu^{\phi}_{,j}(y(t)) f^{\phi}_{,j}(y(t)) + (1 - \mu^{\phi}_{,j}(y(t))) f^{\phi}_{,j+1}(y(t)),$$

where we have introduced the *feedback weights* $\mu^{\phi}_{,j}(x) = \mu_{i,j}(\phi_i(x))$. Thus, the Markov strategy profile describes the dynamics as a stationary function of the current state, as is standard; but we consider Filippov solutions, rather than classical solutions, to the dynamic equation (7).

The standard feedback requirement

$$a_{i,j}(t) = \phi_{i,j}(y(t)) \tag{9}$$

for all $j \in \{1, \ldots, J\}$ and for almost all $t \in \mathcal{T}$ such that $y(t) \in \mathcal{X}_j$ is sufficient to ensure that a trajectory–action (y, a_i) pair has the property that y is a state trajectory of (7).

Definition 9. *Given an initial state x, a full strategy profile ϕ, and a state trajectory y of* (8) *with $y(0) = x$, a* Markovian action schedule $a_i = a^{\phi}_i$ *induced by ϕ is an action schedule such that* (9) *holds almost everywhere. The set of Markovian trajectory-action pairs (y, a^{ϕ}_i) for player i, initial state x, and induced by ϕ is denoted $\mathcal{MA}_{i,x,\phi}$.*

A full strategy profile ϕ and an initial state x uniquely specify the resulting state trajectory, except at pull-pull interfaces $\mathcal{I}_j$ where $f^{\phi}_{,j}(x) \leq 0$ and $f^{\phi}_{,j+1}(x) \geq 0$, with at least one of the inequalities strict. At such points, there are infinitely many trajectories that remain at the interface for an initial time interval of positive length, before moving either to the right or to the left. Even restricting to regular trajectories does not fully eliminate the multiplicity: if both inequalities are strict, there is one regular trajectory that moves immediately to the right, and another that moves immediately to the left.

We now define player i's payoffs and optimal trajectories when restricted to Markovian action schedules.

Definition 10. *Given an initial state $x \in \mathcal{X}$ and a strategy profile $\phi_{-i} \in \mathscr{S}_{-i}$ of the remaining players, the* value of the strategy *ϕ_i to player i is*

$$V_i^{\phi}(x) = \sup_{\mathcal{MA}_{i,x,\phi}} U(y, a_i^{\phi}). \tag{10}$$

Player i's best response is Markovian if, for any initial state x, the player cannot do better than choose a Markovian action schedule:

Definition 11. *A* Markovian best response *by player i to a strategy profile ϕ_{-i} is a strategy ϕ_i such that $V_i(x) = V_i^{\phi}(x)$ for all $x \in \mathcal{X}$.*

Given an initial state x, a strategy profile ϕ_{-i}, and a Markovian best response ϕ_i, a Markovian best response trajectory *for player i is a trajectory y with $y(0) = x$, such that, if a_i^{ϕ} is the action schedule induced by ϕ, the pair (y, a_i^{ϕ}) realises the supremum in* (10)*. The set of Markovian best response trajectories for player i is denoted by $\mathcal{MT}^*_{i,x,\phi}$.*

Finally, we define the game and our equilibrium concept.

Definition 12. *The tuple $\Gamma = (N, \mathcal{X}, \mathcal{Q}_1, \ldots, \mathcal{Q}_N, f, u_1, \ldots, u_N, \rho_1, \ldots, \rho_N)$ defines a* differential game*. If $\mathcal{Q}_i = \mathcal{Q}_j$, $u_i = u_j$ and $\rho_i = \rho_j$ for all i, j, the game is* symmetric.

Definition 13. *A* stationary Markov-perfect Nash equilibrium*, or MPE, of the differential game Γ is a strategy profile $\phi \in \mathscr{S}$ such that, first, for any player i, the strategy ϕ_i is a Markovian best response to ϕ_{-i} and, second, the set of* Markovian equilibrium trajectories *$\mathcal{Y}^*_{x,\phi} = \bigcap_{i=1}^{N} \mathcal{Y}^*_{i,x,\phi_{-i}} \cap \mathcal{MT}^*_{i,x,\phi}$ is nonempty for every $x \in \mathcal{X}$.*

An MPE is continuous *if all ϕ_i are continuous; otherwise it is* discontinuous*. A* symmetric *MPE is an MPE of a symmetric differential game such that $\phi_i = \phi_j$ for all i, j.*

The set of equilibrium trajectories $\mathcal{Y}^*_{x,\phi}$ may contain multiple elements. This gives rise to a problem of trajectory selection, akin to equilibrium selection: different players could choose different trajectories that are consistent with the same strategy profile ϕ. We sidestep this question by assuming that the players are able to coordinate on a joint best response trajectory, which indeed always exists for a symmetric MPE.

5 Results

This section states the main results of our article. The first result shows the well-behavedness of the best-response correspondence for almost all strategy profiles. We start by making precise what we mean "almost all".

Let S be a subset of a complete metric linear space $\mathscr{V}$[16]. The set S is *nowhere dense* if its complement is open and dense; it is *shy* if there exists a Borel set S' containing S and a measure on $\mathscr{V}$ that takes a finite value on some compact set, such that $\text{meas}(S' - v) = 0$ for all $v \in \mathscr{V}$. Shy sets generalise measure zero sets: if $\mathscr{V}$ is finite dimensional, a set is shy if and only if its Lebesgue measure is 0 (Hunt et al., 1992).

Theorem 1. *Let Γ be a differential game for which Assumptions 1, 2, and 3 hold. For every covering $\mathscr{X}$ there is a nowhere dense and shy set $\mathscr{E} \subset \mathscr{S}_{\mathscr{X},-i}$ such that Markovian best response mapping $\mathscr{B}_i : \mathscr{S}_{\mathscr{X},-i} \backslash \mathscr{E} \to \mathscr{S}_i$ is well-defined: for every strategy profile $\phi_{-i} \in \mathscr{S}_{\mathscr{X},-i} \backslash \mathscr{E}$, there is exactly one Markovian best response $\phi_i \in \mathscr{S}_i$ by player i.*

The theorem is proved in Appendices A–D. Appendix A shows that the value function V_i of player i is non-increasing. Appendix B introduces the notion of viscosity solution to the Hamilton–Jacobi–Bellman equation of player i. In Appendix C, we show first that the value function satisfies a number of additional properties: it is left continuous and continuous everywhere except for a finite number of points which we characterise. Then

[16] The space $\mathscr{F}^N_{\mathscr{X}}$ is a complete metric linear space: a metric is constructed as follows. For a compact set $K \subset \mathcal{X}_j$, let $|\phi|_K = \inf\left\{C : \max_{x\in K} |\phi^{(k)}(x)| \leq C^{1+k} k!\right\}$. Let moreover $|\phi|_{j,\infty} = \max_{x\in\mathcal{X}_j} |\phi_j(x)|$ be the max-norm on $\mathcal{X}_j$, where ϕ_j is the extension to $\mathcal{X}_j$ of the restriction of ϕ to $\mathring{\mathcal{X}}_j$. Let $K_{n,j} = [\bar{x}_{j-1} + 1/n, \bar{x}_j - 1/n]$. Introduce the distance

$$d_j(\phi,\psi) = \max\left\{|\phi-\psi|_{j,\infty}, |\phi'-\psi'|_{j,\infty}, \sum_{n=1}^{\infty} 2^{-n} \frac{|\phi-\psi|_{K_{n,j}}}{1+|\phi-\psi|_{K_{n,j}}}\right\}.$$

Then $d(\phi,\psi) = \max_{1\leq j\leq J} d_j(\phi,\psi)$ is a metric on $\mathscr{F}^N_{\mathscr{X}}$, and $\mathscr{F}^N_{\mathscr{X}}$ is complete with respect to d.

we show that the value function is the unique viscosity solution to the Hamilton–Jacobi–Bellman equation in the class of functions with these properties. Appendix D shows that V_i is differentiable on a dense set; using centre manifold theory, we strengthen this to piecewise real analytic. From this the result follows.

The theorem shows that our specification of a differential game, and the Markovian strategy space $\mathscr{S}$, are well-formed in the sense that each player will have a unique best response in $\mathscr{S}_i$ to any profile of the other players' strategies $\mathscr{S}_{\mathscr{X},-i}$ in the complement of the shy set $\mathscr{E}$. Our specification also ensures that payoffs are well-defined for each player, for all strategy profiles $\mathscr{S}$, i.e. the product set of individual strategy spaces—so that no specific admissibility criterion is required.

Another implication is that, while we had to set up in Section 4 the technical apparatus for dealing with potentially non-Markovian best responses, ultimately the best responses turn out to be Markovian, so that for applications it suffices to rely on the simpler Markovian best responses and the associated Filippov dynamics given by equation (8).

Explicit conditions can be formulated for a strategy profile ϕ_{-i} to be in the domain $\mathscr{S}_{\mathscr{X},-i}\backslash\mathscr{E}$ of the best response map: these are of evident importance for applying Theorem 1. One such condition is given in Appendix D as Corollary 1.

As the best response is piecewise analytic, it can be characterised by classical conditions in the regions of analyticity, and by compatibility conditions at the interfaces. The second main result of this article, Theorem 2, formulates such conditions in terms of the value V_i^{ϕ} of strategy ϕ_i to player i, which is the payoff to player i if the strategy profile ϕ is played. Similarly $f_j^{\phi}(x)$ are the local dynamics under the profile ϕ.

Theorem 2. *Assume the same conditions hold as for Theorem 1, and let $\mathscr{E}$ be the shy sets given by that theorem. Let $\phi \in \mathscr{S}$ be such that, for $\mathscr{X}$ a covering adapted to ϕ_{-i}, we have $\phi_{-i} \in \mathscr{S}_{\mathscr{X},-i}\backslash\mathscr{E}$.*

Then $\phi_i = \mathscr{B}(\phi_{-i})$ if and only if the following hold.

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

$$q_i \mapsto u_i(x, q_i) + (V_i^{\phi})'(x) f_{i,j}(x, q_i) \quad \text{on } \mathfrak{Q}.$$

(ii) Monotonicity: V_i^{ϕ} is decreasing and left-continuous.

(iii) Boundary values: If $x = \bar{x}_0$, either $V_i^\phi(x) \geq \beta_i(x)$ or $f_{i,1}(x, q) \geq 0$ for all $q \in \mathcal{Q}$; if $x = \bar{x}_J$, either $V_i^\phi(x) \geq \beta_i(x)$ or $f_{i,J}(x, q) \leq 0$ for all $q \in \mathcal{Q}$.

(iv) Value discontinuities: If V_i^ϕ is not continuous at x, then $x = \bar{x}_j \in \mathcal{I}$, $f_j^\phi(x) \leq 0$ and $f_{i,j+1}(x, q) \geq 0$ for all $q \in \mathcal{Q}$.

(v) Value at interface steady states: For $x = \bar{x}_j \in \mathcal{I}$, let

$$C_{0,j} = \{(q_{i,j}, q_{i,j+1}) \,:\, \mu_{i,j}(q_i) f_{i,j}(x, q_{i,j}) + (1 - \mu_{i,j}(q_i)) f_{i,j+1}(x, q_{i,j+1}) = 0\}\,.$$

Then $\rho V_i^\phi(x) \geq \max_{C_{0,j}} [\mu_{i,j}(q_i) u_i(x, q_{i,j}) + (1 - \mu_{i,j}(q_i)) u_i(x, q_{i,j+1})]$.

(vi) Regularity at strong push–push steady states: If $x = \bar{x}_j \in \mathcal{I}$ is such that

$$\lim_{z \uparrow x} f_{i,j}^\phi(z) > 0 > \lim_{z \downarrow x} f_{i,j+1}^\phi(z),$$

then V_i^ϕ is differentiable at x.

The conditions can be interpreted. Condition (i) is standard. 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. Suppose this were the case and, for intuition, consider flow felicity functions without a bliss point in terms of the control variable. Player i would set the maximal emission rate to grow the stock as fast as possible, at least until the value peaks. But then their flow utility will have been decreasing, as emissions have been constant but damages from the stock have been increasing.

Condition (iii) states that, on the edge of the state space, either a player can exit and take the associated payoff, or exit is impossible.

The remaining conditions place restrictions on the best response where the other players' dynamics are discontinuous. Condition (iv) says that a discontinuity in value is only possible at points where at least one of the other players' strategies is discontinuous, in such a way 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. Condition (ii) then implies the value can only have a downward (not upward) discontinuity at such a point.

Condition (v) ensures that the value at an interface point is at least the value that can be obtained by stabilising the dynamics at that point. Finally, Condition (vi) follows from the fact that value is continuous. If a player's best response is to be pushed strictly towards a stabilisation point, they end up at the same point whether approaching from the left or the right, and very close to the stabilisation point the continuity of the payoffs implies that the marginal value of the stock does not depend on the direction of approach. This implies that, even though our specification allows for situations in which flow payoffs are computed as a convex combination of the payoff for two different actions—which could be interpreted as actions switching "infinitely often"—this outcome never arises in equilibrium.

The proof of Theorem 2 again uses viscosity theory, and is detailed in Appendix E. It shows that the HJB equation of player i has the player's value function as unique viscosity solution. The necessary conditions placed on player i's strategy ϕ_i ensure that the function V_i^ϕ is a viscosity solution of the HJB equation, and therefore equals the value function V_i. This then establishes that ϕ_i is a Markovian best response.

6 Application

We now illustrate how Theorems 1 and 2 can be used in applications. To do this, we consider the canonical transboundary stock pollution game of van der Ploeg and de Zeeuw (1992), introduced in Section 2. The linear–quadratic specification ensures an MPE in piecewise linear strategies that are defined on the whole state space. Dockner and 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).[17] However, these strategies are not supported on the entire state space; as the outcome is not defined in regions of the state space which could be reached with alternative strategies, it is not clear in what sense the purported equilibrium strategies can be said to be best responses (Rowat, 2007).

We shall use Theorem 2 to construct any symmetric equilibrium strategy, as follows. Theorem 2 gives a number of *local* conditions for a best response, valid in the neighbourhood of any x. Any best response is constructed piecewise using continuous elements characterised by Condition (i) of the theorem. The other conditions impose restrictions for how the pieces can be spliced together. Taken together, these conditions determine *global*

[17] Rubio and Casino (2002) show a further class of equilibria. Our methods subsume all these equilibria, showing how they are extended to globally defined discontinuous strategies.

restrictions on any equilibrium strategy, ruling out candidate equilibria which violate one or more necessary conditions. Computation of equilibria is straightforward: after determining a number of important loci, one has to integrate an ordinary differential equation where strategies are continuous, and use the necessary conditions to determine whether any desired discontinuity is allowed, and how the two neighbouring pieces are connected.

The model has been introduced in Section 2. To make the algebra cleaner, we define natural scales of the model, for value $V_0 = \alpha^2/\sqrt{\beta\gamma} = 151$ T\$, time $T_0 = \sqrt{\beta/\gamma} = 417$ y, and pollution $X_0 = \alpha/\sqrt{\beta\gamma} = 0.833$ TtC. Expressing all quantities in multiples of these units leads to a model where $\alpha = \beta = \gamma = 1$, $\rho = 6.25$ and $\delta = 0.417$. For the analytical results we consider N players, and set $\mathcal{Q} = [0, 1]$ and $\mathcal{X} = [0, N/\delta]$.[18]

A symmetric Nash equilibrium strategy $\phi \in \mathscr{S}_i$ for $i = 1, 2, \ldots, N$ and its associated value function V are piecewise real analytic and satisfy at points of differentiability, as a consequence of Condition (i) of Theorem 2, the maximising condition $\phi(x) = q^*(x, V'(x))$, where $q^*(x, p) = p + 1$ if $-1 < p < 0$, $q^*(x, p) = 0$ if $p \le -1$ and $q^*(x, p) = 1$ if $p \ge 0$, as well as the Hamilton–Jacobi equation $\rho V = H(x, V'(x))$, with the *game Hamiltonian* defined by

$$H(x, p) = q^*(x, p) - \frac{1}{2}q^*(x, p)^2 - \frac{1}{2}x^2 + p(Nq^*(x, p) - \delta x).$$

To emphasize, this is not the Hamiltonian to any player's problem, but a Hamiltonian in symmetric equilibrium. The Hamilton–Jacobi equation reads

$$\rho V(x) = \frac{1}{2}(V'(x) + 1)^2 - \frac{1}{2}x^2 + V'(x)\left((N-1)(V'(x) + 1) - \delta x\right). \tag{11}$$

This is a fixed point equation for the value function of the equilibrium strategy that is valid whenever $-1 < V'(x) < 0$. As controls are easier to interpret than shadow values, we use the maximising condition to rewrite the right hand side as

$$\rho V(x) = H_{\text{aux}}(x, \phi(x)) \equiv \frac{1}{2}\phi(x)^2 - \frac{1}{2}x^2 + (\phi(x) - 1)\left((N-1)\phi(x) - \delta x\right). \tag{12}$$

To obtain a differential equation for ϕ, we differentiate with respect to x to obtain

$$\rho V'(x) = \rho(\phi(x) - 1) = \frac{\partial H_{\text{aux}}}{\partial x}(x, \phi(x)) + \frac{\partial H_{\text{aux}}}{\partial q}(x, \phi(x))\phi'(x).$$

[18] I.e. the control ranges between 0 and the bliss point; the state space covers all reachable states.

Solving for ϕ' yields

$$\phi'(x) = \frac{Z_2(x,\phi(x))}{Z_1(x,\phi(x))} = \frac{(\rho+\delta)(\phi(x)-1)+x}{(2N-1)\phi(x)-(N-1)-\delta x}; \tag{13}$$

here $Z_1(x,q)$ and $Z_2(x,q)$ are first degree polynomials in x and q. This equation is a central tool in constructing the set of symmetric Nash equilibria $\phi(x)$ with at most finitely many points of non-differentiability. We plot some integral curves in Figure 4.[19] The pieces out of which we construct any equilibrium must coincide with an integral curve where there are no discontinuities.

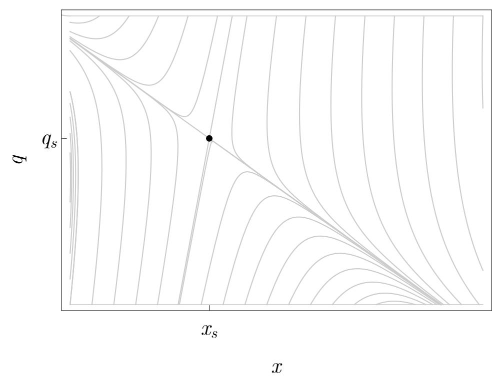

Figure 4: Integral curves of equation (13).

Any piecewise continuously differentiable symmetric Nash equilibrium ϕ has to satisfy equation (13). But as the right hand side is real analytic close to points (x,q) for which $Z_1(x,q) \neq 0$, it follows that in a neighbourhood of a point x such that $Z_1(x,\phi(x)) \neq 0$, the solution ϕ of (13) is real analytic. In particular, a piecewise continuously differentiable equilibrium is necessarily piecewise real analytic.

6.1 Continuous equilibria

We first show that the only continuous equilibrium is the piecewise linear one often studied in the literature. Usually this equilibrium is obtained by algebraic methods (namely, guess-and-verify). We obtain it by geometric arguments, an approach which does not require a linear-quadratic specification.

Denote by $\mathcal{R}(\phi)$ the set of all $x \in \mathcal{X}$ such that $q_\ell < \phi(x) < q_u$. A continuous equilibrium

[19] That is, this is a phase diagram for the auxiliary dynamical system $(x'(s), q'(s)) = (Z_1(x(s), q(s)), Z_2(x(s), q(s)))$.

satisfies (13) for all $x \in \mathcal{R}(\phi)$, and takes either of the values q_ℓ and q_u for $x \notin \mathcal{R}(\phi)$. As the integral curves of (13) bend back on themselves if $Z_1(x,q) = 0$ and $Z_2(x,q) \neq 0$, either $Z_1(x,\phi(x)) \neq 0$ for all $x \in \mathcal{R}(\phi)$ or there are points such that $Z_1(x,\phi(x)) = 0$ and $Z_2(x,\phi(x)) = 0$. Solving these equations shows that there is indeed a unique singularity (x_{ss}, q_{ss}) for which $Z_1(x_{ss}, q_{ss}) = Z_2(x_{ss}, q_{ss}) = 0$.

Restricted to $\mathcal{R}(\phi)$, the differential equation has two linear solutions $w_s(x)$ and $w_u(x)$ that intersect at the singularity, of the form $w_i(x) = q_{ss} + r_i(x - x_{ss})$, $i \in \{s,u\}$. The gradient $r = r_i$ is found by inserting this expression into (13), yielding the condition $(2N-1)r^2 - (\rho + 2\delta)r - 1 = 0$. We find, setting $R = \sqrt{2N - 1 + (\rho/2 + \delta)^2}$, that

$$r_u = \frac{\rho/2 + \delta + R}{2N-1}, \qquad r_s = \frac{\rho/2 + \delta - R}{2N-1}.$$

Given the calibration, all other continuous solutions $\phi(x)$ satisfy either the condition $\phi(x) \geq \max\{w_s(x), w_u(x)\}$ for all $x \in \mathcal{R}(\phi)$, or $\phi(x) \leq \min\{w_s(x), w_u(x)\}$ for all $x \in \mathcal{R}(\phi)$.[20] For those satisfying the former condition, there is a state $x_u < \max \mathcal{X}$ such that $\phi(x) = q_u = 1$ if $x \geq x_u$; in the other case, there is x_ℓ such that $\phi(x) = q_\ell = 0$ if $x \leq x_\ell$. In the first case, for $x \geq x_u$, the value function satisfies $V'(x) \geq q_u - 1 = 0$ for all $x \geq x_u$. As moreover necessarily $V'(x) \leq 0$ for almost all $x \geq 0$ (Theorem 2.ii), this is only possible if $V'(x) = 0$, i.e., if V is constant for all $x \geq x_u$; this is impossible if $x_u < x_{\max}$, as V also has to satisfy

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

In the second case that $\phi(x) \leq \min\{w_s(x), w_u(x)\}$ for all x, the value function satisfies $V'(x_{ss}) = -1$ and $V'(x) \leq q_\ell - 1 = -1$ for $0 \leq x \leq x_\ell$. Moreover, in that interval it satisfies the differential equation

$$\rho V(x) = -\frac{1}{2}x^2 - \delta x V'(x) \quad \text{or} \quad V'(x) = -\frac{\frac{1}{2}x^2 + \rho V(x)}{\delta x}. \tag{14}$$

The graph of a solution V of this differential equation is traced out by integral curves of the dynamical system $x' = -\delta x$ and $v' = \frac{1}{2}x^2 + \rho v$, which has a saddle steady state

[20] For a different parameterisation, a further continuum of continuous equilibria exists if the unstable manifold of system (13) crosses the vertical axis above the origin, in which case all orbits starting on the axis between $w^u(0)$ and the origin are MPE (Rowat, 2007). We omit details for brevity.

at $(\bar{x}, \bar{v}) = (0, 0)$ with unstable eigenspace spanned by $(0, 1)$ and the stable eigenspace spanned by $(1, 0)$. Hence a solution V of (14) is either unbounded as $x \downarrow 0$ or it is differentiable at 0 and $V(0) = V'(0) = 0$. The former situation is ruled out as V is necessarily finite; the latter situation is ruled out as $V'(0) \leq -1$.

The same arguments rule out the possibility that $w_u(x)$ is a Nash equilibrium strategy. Defining $N_2(x)$ by $Z_2(x, N_2(x)) = 0$, we have $w_s(0) < N_2(0) = 1$ and we conclude that the single continuous Nash equilibrium strategy is piecewise linear, and equals $w_s(x)$ for $0 \leq x < x_{s,\ell}$, and 0 for $x \geq x_{s,\ell}$, where $x_{s,\ell}$ is the unique solution of $w_s(x) = q_\ell = 0$.

6.2 Discontinuous continuous-value equilibria

Next we construct all discontinuous equilibria with finitely many discontinuities whose value function is continuous. We first rule out points (x, q) that cannot belong to the graph of any equilibrium strategy (Lemmas 6.1 and 6.2 and Figures 5 and 6). These are either excluded by the arguments used above to rule out strategies which take corner values in continuous equilibria, or by the impossibility of a globally defined, continuous-value strategy passing through them. The remaining points are further divided into regions where there have to be an odd number or an even number of discontinuities or "jumps".[21]

We first define several useful loci in (x, q)-space. At a discontinuity $\bar{x}$ where the value function V is continuous, we write q for one of the limits $\phi_-(\bar{x})$ and $\phi_+(\bar{x})$ of $\phi(x)$ as $x \to \bar{x}$ respectively from below and above, and we introduce $q^\dagger$ for the other one. We then have $H_{\text{aux}}(x, q) = H_{\text{aux}}(x, q^\dagger)$ and $q \neq q^\dagger$, from which we derive the condition $\left(N - \frac{1}{2}\right)(q + q^\dagger) - ((N-1) + \delta x) = 0$, implying

$$q^\dagger = \frac{2(N-1) + 2\delta x}{2N - 1} - q. \tag{15}$$

Note that $q = q^\dagger$ whenever $Z_1(x, q) = 0$. We call two actions q and $q^\dagger$ *conjugate* if they satisfy this relation.

The first pair of useful loci are the *dynamic isocline* D, given by $q = d(x) = \delta x / N$, and

[21] We emphasize that we refer to discontinuities of the policy rule, not along a time trajectory, although a jump in the policy rule could of course produce a discontinuous control trajectory.

the *conjugate dynamic isocline* $D^\dagger$, given by

$$q = d^\dagger(x) = \frac{2(N-1)+2\delta x}{2N-1} - \frac{\delta x}{N}.$$

Recall that an equilibrium strategy induces a state evolution

$$\dot{y}(t) = N\phi(y(t)) - \delta y(t), \qquad y(0) = x.$$

Hence if $(x, \phi(x)) \in D$ then x is a steady state. We let J be the region that is bounded by D and $D^\dagger$ (Figure 5).

The second pair of loci are the graphs of the unique solution ϕ_t of (13) whose graph is tangent to D, and that of the unique solution, denoted ϕ_c, whose graph is tangent to $D^\dagger$, see Figure 5. The first solution has to satisfy $\phi_t(x) = d(x)$ and $\phi_t'(x) = d'(x)$ for some $x = x_t$. This evaluates to

$$\frac{\delta}{N} = \phi_t'(x) = \frac{\rho+\delta}{N-1} + \frac{x}{(N-1)(\delta x/N - 1)},$$

which is solved by

$$x_t = \frac{(\rho + \delta/N)}{(\delta/N)^2 + \rho\delta/N + 1}.$$

For the analogous statement for the conjugate dynamic isocline, we have to solve $\phi_c(x) = d^\dagger(x)$ and $(\phi_c)'(x) = (d^\dagger)'(x)$. This yields

$$\frac{2\delta}{2N-1} - \frac{\delta}{N} = (\phi_c)'(x) = \frac{(\rho+\delta)\left(\frac{2(N-1)+2\delta x}{2N-1} - \frac{\delta x}{N} - 1\right) + x}{(2N-1)\left(\frac{2(N-1)+2\delta x}{2N-1} - \frac{\delta x}{N}\right) - (N-1) - \delta x}$$

which is solved by

$$x_c = \frac{(\rho/(2N-1) + \delta/N)}{(\delta/N)^2 + (\rho/(2N-1))\delta/N + 1}.$$

We note that $x_c < x_t$ and that the graph of any strategy below $\phi_t(x)$ or above $\phi_c(x)$ is disjoint from J.

We now show that a jump in ϕ is only possible if $\dot{x}$ has (weakly) the same sign on both sides of the jump: in the contrary case, we either have a strong push-push or a strong pull-pull point, but this cannot occur in equilibrium.[22]

Proposition 6.1. *A continuous-value Nash equilibrium of the transboundary pollution game has no strong push-push or strong pull-pull points.*

Proof. For the strong push-push case, Theorem 2.vi implies best response value functions are differentiable, and consequently any equilibrium strategy must be continuous. In a Nash equilibrium, all value functions are best responses, and hence all strategies are continuous, yielding a contradiction.

Next, suppose there exists a strong pull-pull point at $x = x_d < x_t$; that is, ϕ satisfies $\phi_-(x_d) < d(x_d)$, $\phi_+(x_d) > d^\dagger(x_d)$. We show that ϕ cannot be extended across the interval $x \in [0, x_d)$. Consider the dynamical system $x' = Z_1(x, q)$, $q' = Z_2(x, q)$, whose orbits coincide with the solutions of (13), and the trajectory with initial point $(x(0), q(0)) = (x_d, \phi_-(x_d))$. For $s > 0$, the trajectory $(x(s), q(s))$ satisfies $q(s) < d(x(s))$ for all $s > 0$: hence there exists $s_0 > 0$ such that $q(s_0) = 0$. By arguments made in the previous section, this orbit cannot represent a ϕ which extends continuously over $x \in [0, x_d]$. Thus there must be a jump at $x_1 \in (0, x_d)$. But as $\phi_+(x_1) < d(x_1)$, value continuity implies $\phi_-(x_1) > d^\dagger(x_1) > d(x_1)$ so that the discontinuity would be strong push-push, which has already been ruled out.

An identical argument, applied to the orbit satisfying $(x(0), q(0)) = (x_d, \phi_+(x_d))$, rules out a strong pull-pull point at $x = x_d > x_c$. As $x_c < x_t$, the argument is complete. □

This argument rules out ϕ satisfying $\phi(x) < \phi_t(x)$, or $\phi(x) > \phi_c(x)$, for any x; as well as $\phi(x) < d(x)$ for $x < x_t$, or $\phi(x) > d^\dagger(x)$ for $x > x_c$, as such strategies can be extended neither continuously nor discontinuously over $\mathcal{X}$. We thus get our first global restriction on regions which a symmetric equilibrium strategy cannot pass through (see Figure 5).

Lemma 6.1. *Let $\phi : \mathcal{X} \to \mathcal{Q}$ be a symmetric continuous-value MPE. Then for any $x \in \mathcal{X}$, $\phi(x) \notin NV_u \cup NV_\ell$, where*

$$NV_u = \big\{(x, q) \,:\, (0 \le x \le x_c \textit{ and } \phi_c(x) < q \le 1) \textit{ or } (x_c < x \le x_{\max} \textit{ and } d^\dagger(x) < q \le 1)\big\}$$

$$NV_\ell = \{(x, q) \,:\, (0 \le x \le x_t \textit{ and } 0 \le q < d(x)) \textit{ or } (x_t < x \le x_{\max} \textit{ and } 0 \le q < \phi_t(x))\}\,.$$

[22] That is, a pull-pull or push-push point in which both inequalities are strict.

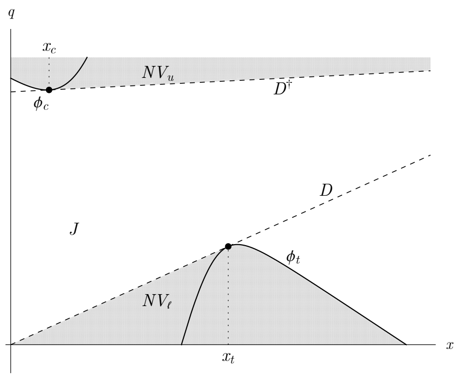

Figure 5: Continuous-value jumps can only occur in the region J bounded by the dynamic isocline D and its conjugate $D^\dagger$. Strategies below ϕ_t, below D for $x < x_t$, above $\phi_t^\dagger$, or above $D^\dagger$ for $x > x_c$ cannot jump, nor be extended continuously to a global equilibrium strategy.

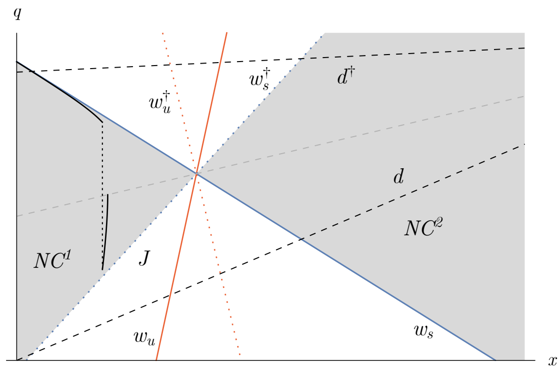

Figure 6: Solutions in the grey region $NC = NC^1 \cup NC^2$ bounded by w_s and $w_u^\dagger$ cannot be continued to globally defined equilibria.

We now turn to points such that equilibria passing through them cannot be continued globally. The situation is illustrated in Figure 6. Let NC denote the region bounded by the graphs of w_s and $w_s^\dagger$: recall that, by equation (15), these graphs cross at (x_{ss}, q_{ss}). Let (x_0, q_0) be an interior point in this region, and assume that ϕ is a globally defined equilibrium that passes through (x_0, q_0). Let $V^\phi(x) = H_{\text{aux}}(x, \phi(x))/\rho$ and $V_s(x) = H_{\text{aux}}(x, w_s(x))/\rho$ be the value functions associated to ϕ and w_s. We can write $NC = \{(x, q) : H_{\text{aux}}(x, q) < \rho V_s(x)\}$.

For $0 \le x < x_{ss}$, the point $(x, \phi(x)) \in NC$ if and only if $w_s^\dagger(x) < \phi(x) < w_s(x)$. This implies if ϕ is differentiable at x that

$$(V^\phi)'(x) = \phi(x) - 1 < w_s(x) - 1 = V_s'(x).$$

At jump points $\bar{x}$, the function V^ϕ is continuous, and both $(\bar{x}, q)$ and $(\bar{x}, q^\dagger)$ are in NC.

It follows that the difference $V_s(x) - V^\phi(x)$ is strictly increasing in x, and therefore

$$V^\phi(x) + (V_s(x_0) - V^\phi(x_0)) < V_s(x)$$

for $x_0 < x \leq x_{ss}$. This implies

$$H_{\mathrm{aux}}(x_{ss}, \phi(x_{ss})) + \rho(V_s(x_0) - V^\phi(x_0)) < H_{\mathrm{aux}}(x_{ss}, q_{ss}) \leq H_{\mathrm{aux}}(x_{ss}, \phi(x_{ss})) :$$

the last inequality holds because $H_{\mathrm{aux}}(x_{ss}, q)$ is minimal at $q = q_{ss}$. We have derived a contradiction, implying that there is no globally defined equilibrium passing through (x_0, q_0). The argument for $x_{ss} < x_0$ is analogous, taking into account that then $V_s(x) - V^\phi(x)$ is strictly decreasing in x.

We thus get our second global restriction on symmetric equilibrium strategies (Figure 6):

Lemma 6.2. *Let $\phi : \mathcal{X} \to \mathcal{Q}$ be a symmetric continuous-value MPE. Then for any $x \in \mathcal{X}$, $\phi(x) \notin NC = NC^1 \cup NC^2$, where*

$$\begin{aligned} NC^1 &= \left\{(x, q)\ :\ w_s^\dagger(x) \leq q < w_s(x)\right\} \\ NC^2 &= \left\{(x, q)\ :\ \max\{0, w_s(x)\} < q \leq w_s^\dagger(s)\right\}. \end{aligned}$$

In fact, by the definition of the set NC, we have shown that:

Theorem 3. *The MPE with the linear strategy w_s achieves the lowest possible payoff compared to all other continuous value MPEs.*

Thus the linear feedback Nash strategy, a focus of the literature since Starr and Ho (1969), is Pareto-dominated by all other strategies associated with a continuous value function.

Why is the linear feedback strategy the worst equilibrium? Because it is an equilibrium in which players respond relatively passively to additional emissions. Consider $x = x_{ss}$; for this state, the linear strategy has the steepest downward slope of any feasible strategy, whether locally or globally defined (see Figure 4). For this strategy, a marginal increase in the carbon stock leads all players to reduce their emissions. In the climate mitigation game, such passive responses are bad, as they give a great incentive for any player to emit more. All other feasible strategies are more aggressive at $x = x_{ss}$; a marginal unit of the stock leads to an increase in emissions. For this state, the linear strategy is the worst possible one, precisely as the players' strategic responses are so weak.

For states $x \neq x_{ss}$, it is feasible to construct even worse local equilibrium strategies, with even weaker responses. However, these cannot be extended globally: such a strategy is Pareto-inferior to the linear strategy all along its extension—but the linear one is the worst one at $x = x_{ss}$, so that a strategy weaker than the linear strategy cannot be extended across this point with a continuous value.[23] The main lesson is thus that equilibria with aggressive responses yield higher welfare for all players.

We can now characterise all Nash equilibria with finitely many discontinuities. The result divides the (x, q) plane in a number of regions.

Let x_1 be the unique solution of $d^\dagger(x) = w_s(x)$ and x_2 the unique solution of $d(x) = w_s(x)$. The allowed regions are:

$$\begin{aligned}
C_u &= \{(x,q) \,:\, (0 \le x < x_1 \text{ and } w_s(x) \le q \le \phi_c(x)) \\
&\qquad\qquad \text{or } (x_1 \le x < x_c \text{ and } d^\dagger(x) < q < \phi_c(x)\} \\
C_\ell &= \{(x,q) \,:\, (x_t < x < x_2 \text{ and } \phi_t(x) \le q < d(x)) \\
&\qquad\qquad \text{or } (x_2 \le x \le x_{\max} \text{ and } \max\{\phi_t(x), 0\} \le q \le \max\{w_s(x), 0\})\} \\
OJ_u &= \left\{(x,q) \,:\, \max\{w_u(x), w_u^\dagger(x)\} \le q \le d^\dagger(x)\right\}, \qquad OJ_\ell = (OJ_u)^\dagger, \\
EJ_u^1 &= \left\{(x,q) \,:\, w_s(x) \le q < \min\{d^\dagger(x), w_u^\dagger(x)\}\right\}, \qquad EJ_\ell^1 = (EJ_u^1)^\dagger \\
EJ_u^2 &= \left\{(x,q) \,:\, w_s^\dagger(x) < q < \min\{d^\dagger(x), w_u(x)\}\right\}, \qquad EJ_\ell^2 = (EJ_u^2)^\dagger.
\end{aligned}$$

The forbidden regions $NC^1, NC^2, NV_u, NV_\ell$ have been defined earlier.

The next result characterises the set of all continuous–value Markov perfect Nash equilibria of the transboundary pollution game with finitely many jump points.

Theorem 4. *A piecewise real analytic function $\phi : \mathcal{X} \to \mathcal{Q}$ with finitely many points of non-differentiability is a continuous value Markov perfect Nash equilibrium of the transboundary pollution game if and only the following conditions hold.*

[23] In Section 6.3, we show that value discontinuities admit equilibria with worse strategies.

(i) Its graph is contained in the set

$$C_u \cup C_\ell \cup OJ_u \cup OJ_\ell \cup EJ_u^1 \cup EJ_\ell^1 \cup EJ_u^2 \cup EJ_\ell^2.$$

(ii) At points x where ϕ is differentiable and $\phi(x) > 0$, equation (13) *holds.*

(iii) At a point of non-differentiability $\bar{x}$, the values $\phi_-(\bar{x}) = \lim_{x\uparrow\bar{x}}$ and $\phi_+(\bar{x}) = \lim_{x\downarrow\bar{x}}$ are conjugate.

Proof. This is a direct consequence of Theorem 2. The value function is continuous by assumption, and the dynamics are inward pointing for all actions on the boundary of $\mathcal{X}$.

The necessity of the conditions has been used in the construction given above.

Sufficiency: equation (13) implies Condition (i) of Theorem 2. Since $\phi(x) < 1$, it follows that $V'(x) < 0$ and V is decreasing, implying Condition (ii). The dynamics are inward pointing at the boundary of $\mathcal{X}$, implying Condition (iii). The function V is continuous by hypothesis, taking care of Condition (iv). As there are no strong push-push or pull-pull steady states, an interface is a steady state of either the left-hand or the right-hand dynamics, which implies Condition (v) and shows that Condition (vi) is void. □

The result is illustrated in Figure 7. Jumps are possible from points in the intersection of J and the complement of NC. Going through the diagram with increasing values of x, a jump is possible from the region marked EJ_u^1 to EJ_ℓ^1. But if such an equilibrium would not jump a second time, it would enter the region NC^1, which is ruled out. Hence there must be an even number of jumps between the regions EJ_u^1 and EJ_ℓ^1, and similarly between EJ_u^2 and EJ_ℓ^2. Jumps are also possible from OJ_u to OJ_ℓ. In this situation, an odd number of jumps is required, because otherwise the equilibrium has to enter NV_u, which is again ruled out.

Note that the Hamiltonian is convex in q and takes a minimum where $Z_1(x, \phi(x)) = 0$ (see equation (13), illustrated in Figure 7 by the grey dashed line). Thus the value $V^\phi(x)$ of an equilibrium strategy ϕ at any x increases away from the grey dashed line, and the Pareto-dominant strategy for any x can be obtained with a strategy coinciding, at x, with the top or bottom boundary of the allowable region.

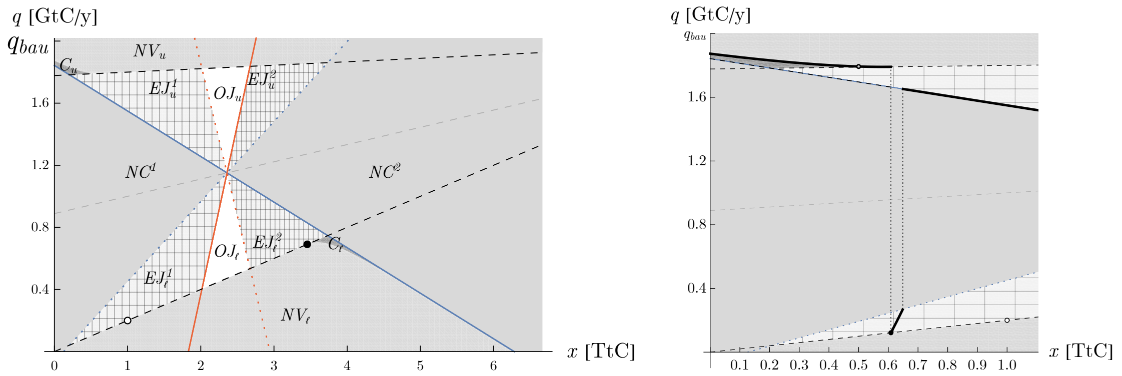

(a) Regions

(b) Pareto-optimal strategy

Figure 7: Feedback Nash equilibria cannot enter the region $NC^1 \cup NC^2 \cup NV_u \cup NV_\ell$ (dark grey). Equilibria have no discontinuities in the region $C_u \cup C_\ell$ (light grey), an odd number of discontinuities in the region $OJ_u \cup OJ_\ell$ (white), and an even number of discontinuities in the region $EJ_u^1 \cup EJ_\ell^1 \cup EJ_u^2 \cup EJ_\ell^2$ (crossed). The equilibrium strategy of Figure 1(a) that is Pareto-optimal for the initial state $x = 0.5$ TtC is an example of a strategy with two jumps.

In practice, any continuous-value equilibrium ϕ can be computed in the following manner. Pick a desired point $(x_0, \phi(x_0))$. Extend the strategy to the left along the integral curve of equation (13). Insert a discontinuity at desired point x; the discontinuity must connect $\phi(x)$ to $\phi(x)^\dagger$. A discontinuity must be inserted if the extension meets a boundary of the allowable region. Continue extending and inserting discontinuities until ϕ extends to $x = 0$. Then do the same to the right of x_0. This is a computationally simple exercise, especially once the desired loci have been computed, as it simply requires computation of some integral curves.

The Pareto-dominant equilibrium, for the calibrated initial state, has been given in Figure 1(a); see Section 2 for discussion of the intuition. Note that the steady state carbon concentration is below the socially optimal long-run stock. More generally, the Pareto-dominant equilibrium is conditional on the initial state: that is, an equilibrium which may be dominant at the initial state may reach a state at which a different equilibrium dominates it in terms of the continuation payoff.[24] The envelope of values achievable in a symmetric equilibrium over and above the business-as-usual (non-equilibrium) value has been given in Figure 2. Discontinuous equilibria can do substantially better than the linear equilibrium; with our calibration, the linear equilibrium closes just over 50% of the difference between business-as-usual and the social optimum, whereas choosing the best

[24]This outcome highlights the fact that MPE may not be renegotiation-proof. We leave exploration of renegotiation for future work.

symmetric equilibrium can close between 50-100% of the remaining difference, and can indeed sustain the socially optimal steady state.[25]

6.3 Discontinuous-value discontinuous equilibria

In addition to continuous-value equilibria, we now discuss the possibility of equilibria involving a discontinuity in the value function. Suppose that there is an equilibrium strategy $\phi(x)$ such that the value is discontinuous at $x = \bar{\bar{x}}$. Parts (ii) and (iv) of Theorem 2 then imply, first, that $V_-(\bar{\bar{x}}) \geq V^{\daleth}(\bar{\bar{x}}) > V_+(\bar{\bar{x}})$, and, second, that the equilibrium strategy satisfies $\phi_-(\bar{\bar{x}}) \leq d(\bar{\bar{x}})$ and $\phi_+(\bar{\bar{x}}) > \frac{\delta \bar{\bar{x}}}{N-1} \equiv \ell^{\mathrm{r}}(\bar{\bar{x}})$, where we introduce the *right non-controllability locus* as the graph of the function $\ell^{\mathrm{r}}(x)$.

Assume that $\phi \notin \mathscr{E}$: then Theorem 2 implies that the number of discontinuities is finite, and hence also the number of discontinuous-value discontinuities. Let $\bar{\bar{x}}$ denote the smallest of these. We cannot have that $\bar{\bar{x}} \leq x_t$, for $\phi(x)$ has to be continuous in an interval $(\bar{\bar{x}} - \delta, \bar{\bar{x}})$ for some $\delta > 0$. But then we obtain for x in this interval that $\phi(x) < d(x)$. The value associated to ϕ is continuous in $(0, \bar{\bar{x}})$; but such an equilibrium strategy has been ruled out by the results of the previous subsection.

Hence $\bar{\bar{x}} > x_t$. We have necessarily $\phi_-(\bar{\bar{x}}) \leq d(x)$, because of Theorem 2.(iv), as well as $\phi_+(\bar{\bar{x}}) < (\phi_-(\bar{\bar{x}}))^\dagger$, to ensure that $V_-(\bar{\bar{x}}) > V_+(\bar{\bar{x}})$, and $\phi_+(\bar{\bar{x}}) < d^\dagger(\bar{\bar{x}})$, otherwise ϕ cannot be extended for $x > \bar{\bar{x}}$ (again by the arguments of the previous subsection). In fact, as $\phi_-(\bar{\bar{x}}) \leq d(\bar{\bar{x}})$, the second inequality implies the first, and we are left with the condition

$$\ell^{\mathrm{r}}(\bar{\bar{x}}) \leq \phi_+(\bar{\bar{x}}) < d^\dagger(\bar{\bar{x}}).$$

In particular, as $\ell^{\mathrm{r}}(x) = d^\dagger(x)$ for $x = x^o \equiv (N-1)^2/((N-1/2)^2 + 1/4)(N/\delta)$, we find the necessary and sufficient condition $x_t < x^o$ for the existence of discontinuous-value discontinuities, which can be expressed as

$$\rho < \frac{2N^2(N-1)^2 - (2N-1)\delta}{N(2N-1)\delta}.$$

A discontinuous value jump starts from $VD_- = \{(x, q) : x > x_t \text{ and } \phi_t(x) \leq q < d(x)\}$

[25] More complicated equilibrium strategies can achieve an arbitrarily close approximation of the value envelope (cf. Schumacher et al., 2022).

and ends in $VD_+ = \{(x,q) : x > x_t \text{ and } \ell_r(x) \leq q < d^\dagger(x)\}$. More precisely, the smallest discontinuous value jump has to be from the region $VD_{-,1} = VD_- \backslash NC^2$, as otherwise it cannot be continuous for smaller values of x. If it lands in $VD_{+,1} = VD_+ \backslash NC^2$, it has for some larger value of the state to jump down with a continuous value jump back to EJ_ℓ^2; if in $VD_{+,2} = VD_+ \cap \{(x,q) : Z_1(x,q) > 0\}$, then it has to jump down with a continuous value jump to $\{(x,q) : q \geq d(x)\}$; if it finally lands in $VD_{+,3} = VD_+ \cap \{(x,q) : Z_1(x,q) < 0\}$, then it can be continued continuously.

Once the strategy enters for stock values larger than the smallest discontinuous value jump again the region VD_-, it can make another discontinuous value jump. These jumps are illustrated in Figure 8.

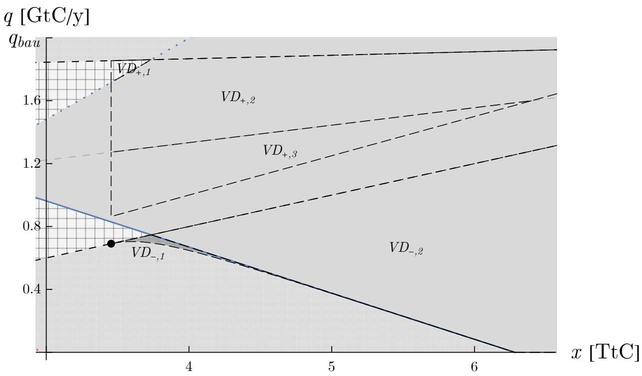

(a) Regions

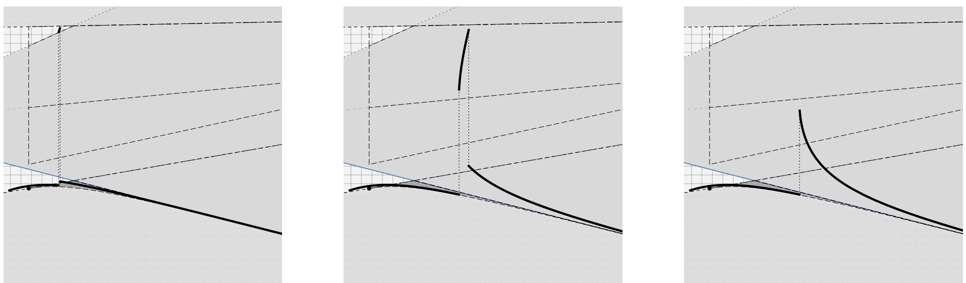

(b) Jump to $VD_{+,1}$ (c) Jump to $VD_{+,2}$ (d) Jump to $VD_{+,3}$

Figure 8: The value discontinuity with the smallest stock value can only occur as upward jumps from the region $VD_{-,1}$ to the region $VD_+ = VD_{+,1} \cup VD_{+,2} \cup VD_{+,3}$.

The intuition is that if the initial level of pollution is high, then it is possible to get stuck in 'strategic trap' equilibria. The 'trap' is a pollution stock threshold such that, if above but very close to the threshold, all players would prefer to cross it. If they could do it, the players would jointly choose lower emission rates and the system would transit to a relatively benign steady state. However, given the high-pollution strategies of the other players, no player is able to individually steer the system below the threshold. Each player best responds with a high pollution rate, leading the system to a high-pollution steady state. This outcome is more likely if the number or players is large, so that each player is unable to reverse the evolution of the pollution stock on their own.[26]

7 Conclusion

We have shown how to put differential games of public investment on a solid conceptual footing, solving technical issues which have bedevilled the analysis of MPE of such games. Introducing a strategy space that includes discontinuous strategies, we have shown that best responses almost always exist and are unique. We have also given a necessary and sufficient condition for constructing the best response.

This condition is straightforward to use in applications, as we have demonstrated with reference to a classic model of non-cooperative mitigation of climate change. We have proved the existence of a large class of symmetric equilibria, all but one of them featuring discontinuous strategies. In other applications, with non-convex natural dynamics, continuous equilibria do not exist at all; our formulation allows meaningful study of MPE in such games also (see Kossioris et al., 2008; Dockner and Wagener, 2014).

The importance of multiple equilibria in our application points to the question of equilibrium selection. Dockner and Long (1993) considered an important function of climate negotiations to be coordination to a relatively benign equilibrium.[27] Our substantive message is twofold. First, in a non-cooperative world, gradual stabilisation of emissions implies poor strategic incentives and thus leads to worse outcomes; in our simple model, the unique equilibrium with gradual stabilisation is the worst possible one. Second, benign long-run steady states are sustained in equilibrium by 'trigger'-like (but Markovian) asymmetric responses, in which deviations from the 'good' steady state are punished by the players choosing to collectively slowly drift to a 'bad' steady state with a higher carbon concentration. The good steady state is reached in finite time, with a discontinuous reduction in carbon emissions, suggesting that an agreeing on radical climate policies

[26] Benhabib and Radner (1992) support a high renewable resource steady state by what is effectively a 'trap', should the stock cross below the steady state. In their case, the ability of the players to consume a large quantity of stock in a discrete time period ensures loss of the ability to drive the stock back up above the trap threshold.

[27] 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).

might be better than agreeing on gradual reductions, specifically because radical policies provide better strategic incentives (similar to Dutta and Sundaram, 1993).[28] Such equilibria can be indexed by a ‘stabilisation target’, followed by ‘net zero’ emissions.

A criticism of our approach is that discontinuous strategies, especially discontinuous control trajectories, are unrealistic. We defend our approach with three arguments. First, real-world policies not only can, but almost always are, adjusted in discrete increments. Carbon taxes, nominal interest rates, or investment rates are in the real world adjusted as a step-wise process of non-infinitesimal changes at specific points in time. Our equilibria can be implemented by simply conditioning e.g. a carbon tax on the atmospheric carbon concentration. If such a policy were to lead to a large change, overnight, in the tax rate, we have little doubt that emissions would respond rather discontinuously.

Second: discontinuous, large changes in emissions may seem implausible. However, even policies with very large impacts are sometimes implemented in a discontinuous manner: consider the abrupt lockdowns in 2020 as SARS-CoV-2 surged around the world.[29] To the extent that such changes seem less plausible in the context of climate policy, this likely reflects a concern over adjustment costs, absent in our model. We argue the appropriate response is to develop models which allow the explicit treatment of such costs. This would require moving to a multidimensional state space; we expect discontinuous strategies to be important also in that context. Our work is a necessary first step towards this.

Third, our simple climate policy example is not intended as a quantitative guide to policy. Instead, it provides qualitative lessons about the structuring of climate agreements, and how strategic leverage can be used to develop agreements with more effective incentives. Our application should not be taken literally, but should be taken seriously.[30]

Our methodological results open up research avenues in numerous directions. The present model is deterministic; which of the equilibria are limits of corresponding stochastic equilibria as noise in the state evolution tends to zero is an open question. We expect that our results can be extended to the case of multiple state variables. We have considered strategies with a finite number of discontinuities. The exceptional set for which the best-

[28] Van Der Ploeg and Venables (2022) give further reasons for radical climate policies, based on strategic complementarities and multiple steady states.

[29] There are also many applications in which large policy changes are unproblematic; consider e.g. financial trading, or interest rate changes.

[30] Rubinstein (2006) argues that economic theory is only loosely connected to reality, but can still produce results which are useful in the real world.

response correspondence does not map back into itself is an irreducible feature, not a bug; it would be interesting to understand whether the correspondence would be exactly closed if countably many discontinuities are allowed. Lastly, we have only considered symmetric equilibria, but our results also allow the study of asymmetric equilibria.

Finally, given that our methods have found several novel and important features of a model which had been fairly extensively studied in the past, we expect these methods to yield new insights in the numerous applications of differential games which exist in the literature. While we have focused on deterministic differential games of public investment, we expect our methods can also facilitate the analysis of various types of games of incomplete information in continuous time.

Acknowledgements

We thank Buz Brock, 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 ifo Institute, the Swedish University of Agricultural Sciences, University of Amsterdam, and University of Bologna, for helpful comments.

References

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

Achdou, Y., Han, J., Lasry, J.M., Lions, P.L., Moll, B., 2022. Income and wealth distribution in macroeconomics: A continuous-time approach. The Review of Economic Studies 89, 45–86.

Aubin, J.P., 1993. Optima and equilibria. Springer, Berlin.

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

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

Barilla, C., Gonçalves, D., 2024. The dynamics of instability. Theoretical Economics 19, 365–405.

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

Barles, G., Briani, A., Chasseigne, E., 2014. A Bellman Approach for Regional Optimal Control Problems in $\mathbb{R}^N$. SIAM Journal on Control and Optimization 52, 1712–1744.

Başar, T., Olsder, G., 1982. Dynamic Noncooperative Game Theory. first ed., SIAM, Philadelphia, PA.

Başar, T., Zaccour, G., 2018. Handbook of Dynamic Game Theory. Springer, Cham.

Benhabib, J., Radner, R., 1992. The joint exploitation of a productive asset: a game-theoretic approach. Economic Theory 2, 155–190.

Board, S., Meyer-ter Vehn, M., 2013. Reputation for quality. Econometrica 81, 2381–2462.

Brunnermeier, M., Sannikov, Y., 2016. Macro, money, and finance: A continuous-time approach, in: Handbook of Macroeconomics. Elsevier. volume 2, pp. 1497–1545.

Cannarsa, P., Frankowska, H., 1991. Some characterizations of optimal trajectories in control theory. SIAM Journal on Control and Optimization 29, 1322–1347.

Cannarsa, P., Frankowska, H., Scarinci, T., 2015. Second-order sensitivity relations and regularity of the value function for Mayer's problem in optimal control. SIAM Journal on Control and Optimization 53, 3642–3672.

Cesari, L., 1983. Optimization — Theory and Applications. Springer, New York.

Cisternas, G., 2018. Two-sided learning and the ratchet principle. The Review of Economic Studies 85, 307–351.

Clarke, F., Ledyaev, Y., Stern, R., Wolenski, P., 1998. Nonsmooth analysis and control theory. Springer, New York.

De Frutos, J., Martín-Herrán, G., 2018. Selection of a Markov perfect Nash equilibrium in a class of differential games. Dynamic Games and Applications 8, 620–636.

Dockner, E., Jørgensen, S., van Long, N., Sorger, G., 2000. Differential Games in Economics and Management Science. CUP, Cambridge, UK.

Dockner, E., Long, N.V., 1993. International pollution control: Cooperative versus noncooperative strategies. Journal of Environmental Economics and Management 25, 13–29.

Dockner, E., Sorger, G., 1996. Existence and properties of equilibria for a dynamic game on productivity assets. Journal of Economic Theory 71, 209–227.

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

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

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

Faingold, E., Sannikov, Y., 2011. Reputation in continuous-time games. Econometrica 79, 773–876.

Fudenberg, D., Tirole, J., 1991. Game Theory. MIT Press, Cambridge, MA.

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

Hirsch, M., Pugh, C., Shub, M., 1977. Invariant manifolds. volume 583 of *Lecture Notes in Mathematics*. Springer, Heidelberg.

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

Keller, G., Rady, S., 2015. Breakdowns. Theoretical Economics 10, 175–202.

Keller, G., Rady, S., Cripps, M., 2005. Strategic experimentation with exponential bandits. Econometrica 73, 39–68.

Klein, N., Rady, S., 2011. Negatively correlated bandits. The Review of Economic Studies 78, 693–732.

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

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

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

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

Rubinstein, A., 2006. Dilemmas of an economic theorist. Econometrica 74, 865–883.

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

Sannikov, Y., 2007. Games with imperfectly observable actions in continuous time. Econometrica 75, 1285–1329.

Sannikov, Y., 2008. A continuous-time version of the principal-agent problem. The Review of Economic Studies 75, 957–984.

Schumacher, J., Reddy, P., Engwerda, J., 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, 79–100.

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

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

Takens, F., Vanderbauwhede, A., 2010. Local invariant manifolds and normal forms, in: Broer, H., Hasselblatt, B., Takens, F. (Eds.), Handbook of Dynamical Systems, Volume 3. North-Holland, Amsterdam. chapter 3, pp. 89–124.

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

Van Der Ploeg, F., Venables, A.J., 2022. Radical climate policies. World Bank Policy Research Working paper 10212.

Vinter, R., 2000. Optimal control. Birkhäuser, Boston.

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

Online appendix

As the remainder of the article treats the dynamic optimisation problem of a single player, given the strategic choices of the other players, we drop the index i, and we replace the index '$, j$', which indicates restriction to the interval $\mathcal{X}_j$, by j. This convention holds unless we are explicitly referring to the game. Also we assume throughout that Assumptions 1, 2, and 3 hold, without explicitly mentioning the fact.

A Singular and regular value functions are identical

In this section, we show that the value function is decreasing.

A.1 Notation

If $S \subset \mathbb{R}$ is a measurable subset, $|S|$ denotes the Lebesgue measure of S. For bounded continuous functions $h : S \to \mathbb{R}$, we set $\|h\|_\infty = \sup_{z \in S} |h(z)|$.

We have that the local dynamics f_j and the local payoffs u_j are bounded, real analytic on $\mathring{\mathcal{X}}_j$, and can be continuously extended, together with their derivatives, to the closed interval $\mathcal{X}_j$. The global dynamics $\mathbf{f}$ and payoffs $\mathbf{u}$ equal their local counterparts f_j and u_j on $\mathring{\mathcal{X}}_j$ and appropriately weighted convex combinations of f_j and f_{j+1}, respectively u_j and u_{j+1}, on interfaces $\mathcal{I}_j$. We therefore have $\|\mathbf{f}\|_\infty = \|f\|_\infty$ and $\|\mathbf{u}\|_\infty = \|u\|_\infty$.

A.2 Existence of optimal action schedules

Proposition A.1. *For every $x \in \mathcal{X}$, there is a trajectory–action pair (y^*, a^*) such that $y^*(0) = x$ and $V(x) = U(y^*, a^*)$.*

The proof of this proposition is given in Appendix F.1; it follows from Assumption 2.

A.3 Equality of singular and regular value

A sufficient condition for uniqueness of solutions to the Hamilton–Jacobi–Bellman equation is equality of value and regular value function (Barles et al., 2013, 2014).

The object of this section is to show this equality for our context. Our first result is a consequence of the assumption $u_x < 0$: given a fixed emission strategy, payoffs decrease with increasing initial pollution stock levels.

Proposition A.2. *Let (y, a) and $(\tilde{y}, a)$ be trajectory–action pairs with the same action schedule and initial points $x, \tilde{x} \in \mathfrak{X}$, $\tilde{x} < x$. If $\tilde{y}(t) \leq y(t)$ for all $t \in \mathfrak{T}$, then $U(\tilde{y}, a) > U(y, a)$.*

Proof. This follows from the facts that $\tilde{y}(t) < y(t)$ in a neighbourhood of $t = 0$ and that $\mathbf{u}(x, q, \lambda)$ is decreasing in $x \in \mathfrak{X}$. $\square$

This idea is used to show that the value function decreases; details are given in Section F.2.

Proposition A.3. *The value function is decreasing.*

To prove equality of value and regular value function, we are going to exhibit for every trajectory–action pair a regular trajectory–action pair generating an outcome that is at least as good. First we show that such a regular trajectory either almost never is at its initial state, or it is there always.

Proposition A.4. *Let (y, a) be a trajectory–action pair with initial state x. Then there is a trajectory–action pair $(\tilde{y}, \tilde{a})$ such that $U(\tilde{y}, \tilde{a}) \geq U(y, a)$ and either $\tilde{y}(t) > x$ for all $t \in \mathfrak{T}$, or $\tilde{y}(t) < x$ for all $t \in \mathfrak{T}$, or $\tilde{y}(t) = x$ and $\tilde{a}(t)$ constant for all $t \in \mathfrak{T}$.*

The proof is given in Section F.3.

If a trajectory–action pair is always at an interface steady state where the regularity condition (6) is not satisfied, that is, at a 'pull-pull' steady state in the terminology of Barles et al. (2013), there is a second trajectory–action pair with the same action schedule and the same initial condition such that the trajectory is always to the left of that steady state, and such that the pair has a higher total payoff. This is the heart of the following result.

Proposition A.5. *For every non-regular trajectory–action pair there is a regular pair with a higher total payoff.*

The proof consists in replacing all singular pull–pull trajectory segments by regular trajectories going to the left. The details are given in Section F.4.

Proposition A.6. *The value function V and the regular value function V^{reg} are equal.*

Proof. Take $x \in \mathcal{X}$ and find a trajectory–action pair (y, a) such that $y(0) = x$ and $V(x) = U(y, a)$. By Proposition A.5 there is a regular trajectory–action pair $(\tilde{y}, \tilde{a})$ such that $\tilde{y}(0) = x$ and $U(\tilde{y}, \tilde{a}) \geq U(y, a)$. We have that $V(x) = U(y, a) \leq U(\tilde{y}, \tilde{a}) \leq V^{\text{reg}}(x) \leq V(x)$, and hence $V(x) = V^{\text{reg}}(x)$. $\square$

B Viscosity solutions

The value function of the optimisation problem is the unique viscosity solution of a Hamilton–Jacobi–Bellman (HJB) equation. We introduce the appropriate notions.

For $j \in \{1, \ldots, J\}$, $x \in \mathcal{X}_j$ and $p \in \mathbb{R}$, the *local Hamilton function* is

$$H_j(x, p) = \max_{q \in \mathcal{Q}} \left[u(x, q) + p f_j(x, q) \right].$$

Assumption 2 implies that the function $q \mapsto u(x, q) + p f_j(x, q)$ has a unique maximiser $q^* = q^*_j(x, p)$ in $\mathcal{Q}$. If $q^* \in \mathring{\mathcal{Q}}$, then $u_q(x, q^*) + p(f_j)_q(x, q^*) = 0$.

We also define local Hamilton functions H_0 and H_{J+1} at the boundary. For this, we introduce functions $f_0(x, q)$, $f_{J+1}(x, q)$ and $n(x)$ as follows. If $x = x_{\min}$, we set $f_0(x, q) = f_1(x, q)$ and $n(x) = -1$; if $x = x_{\max}$, we set $f_{J+1}(x, q) = f_J(x, q)$ and $n(x) = 1$. If for $x \in \partial\mathcal{X}$ and $j \in \{0, J+1\}$ we have $n(x) f_j(x, q) > 0$ for some $q \in \mathcal{Q}$, then exit from $\mathcal{X}$ is possible at x and we set $H_j(x, p) = \rho\beta(x)$. If exit is not possible at x, we set $H_j(x, p) = -\infty$.

The *Hamilton function* of the optimisation problem is

$$\mathbf{H}(x, p) = \begin{cases} H_j(x, p) & x \in \mathring{\mathcal{X}}_j,\ j \in \{1, \ldots, J\}, \\ \max\{H_j(x, p), H_{j+1}(x, p)\} & x \in \mathcal{I}_j,\ j \in \{0, \ldots, J\}. \end{cases}$$

For an interface point $\bar{x}_j \in \mathcal{I}$, denote the set of actions stabilising it by

$$C_{0,j} = \{q \,:\, \mathbf{f}(\bar{x}_j, q) = 0\}\,.$$

The *interface Hamilton function* is then given as

$$H_j^{\mathcal{I}}(\bar{x}_j) = \max_{q \in C_{0,j}} [\mathbf{u}(\bar{x}_j, q)]\,.$$

We set $H_j^{\mathcal{I}}(\bar{x}_j) = -\infty$ if the set $C_{0,j}$ is empty.

Let $\mathcal{Z} \subset \mathbb{R}$. For a function $W : \mathcal{Z} \to \mathbb{R}$, the *upper semi-continuous envelope* is

$$W^*(x) = \lim_{\delta \downarrow 0} \sup\{W(z) \,:\, z \in \mathcal{Z},\, |z - x| \leq \delta\}.$$

The *lower semi-continuous envelope* W_* is defined analogously, with inf replacing sup. We have that $\mathbf{H}^*(x,p) = \mathbf{H}(x,p)$ for all (x,p), while

$$\mathbf{H}_*(x,p) = \begin{cases} H_j(x,p) & x \in \mathring{\mathcal{X}}_j,\ j \in \{1, \ldots, J\}, \\ \min\{H_j(x,p), H_{j+1}(x,p)\} & x \in \mathcal{I}_j,\ j \in \{0, \ldots, J\}. \end{cases}$$

The *superdifferential* $D^+W(x)$ of a bounded upper semicontinuous function $W : \mathcal{Z} \to \mathbb{R}$ at a point x is the set

$$D^+W(x) = \left\{ p \in \mathbb{R} \,:\, \limsup_{\substack{z \to x \\ z \in \mathcal{Z}}} \frac{W(z) - W(x) - p(z-x)}{|x - z|} \leq 0 \right\}.$$

The *subdifferential* $D^-W(x)$ of a bounded lower semicontinuous function W on $\mathcal{Z}$ at x is defined similarly, with sup replaced by inf and $\leq$ by $\geq$ (see e.g. Bardi and Capuzzo-Dolcetta, 2008, Chapter V). We have that $p \in D^+W(x)$ if and only if there is a continuously differentiable function ψ such that $\psi'(x) = p$ and $W - \psi$ restricted to $\mathcal{X}$ has a local maximum at x. An analogous characterisation exists for subdifferentials.

Definition 14. *The function $W : \mathcal{X} \to \mathbb{R}$ is a* viscosity supersolution *of the Hamilton–Jacobi–Bellman equation*

$$\rho W(x) - \mathbf{H}(x, W'(x)) = 0 \tag{16}$$

if for all $x \in \mathcal{X}$ *and all* $p \in D^{-}W_{*}(x)$ *we have*

$$\rho W_{*}(x) - \mathbf{H}_{*}(x, p) \geq 0, \tag{17}$$

and if for all $j \in \{1, \ldots, J-1\}$ *and* $x \in \mathcal{I}_j$*, we have*

$$\rho W_{*}(x) - H^{\mathcal{I}}(x) \geq 0. \tag{18}$$

The function W *is a* viscosity subsolution *of* (16) *if for all* $x \in \mathcal{X}$ *and all* $p \in D^{+}W^{*}(x)$

$$\rho W^{*}(x) - \mathbf{H}^{*}(x, p) \leq 0. \tag{19}$$

Finally, W *is a* viscosity solution *of* (16) *if it is both a supersolution and a subsolution.*

Note that (18) is only a condition for being a supersolution, not a subsolution.

Theorem B.1. *The value function* V *is a viscosity solution of the Hamilton–Jacobi–Bellman equation* (16)*.*

Proof. The statement is local, and the proof is a combination of known results. See Proposition III.2.8 of Bardi and Capuzzo-Dolcetta (2008) for (17) and (19) if $x \in \mathring{\mathcal{X}}_j$; Theorem 2.5 of Barles et al. (2013) for the case that $x \in \mathcal{I}_j$, $j \in \{1, \ldots, J-1\}$, as well as for (18); and Theorem V.4.13 of Bardi and Capuzzo-Dolcetta (2008) for the boundary case $x \in \mathcal{I}_j$ if $j \in \{0, J\}$. $\square$

C Unicity of solutions

A central question is whether the dynamics are controllable at an interface.

Definition 15. *The dynamics are* left (right) controllable *at a point* $\bar{x}_j \in \mathcal{I} \cup \partial X$*, if there is an open interval* I *containing* 0 *such that*

$$I \subset f_j(\bar{x}_j, \mathcal{Q}) \qquad \Big(I \subset f_{j+1}(\bar{x}_j, \mathcal{Q})\Big).$$

If the dynamics are both left and right controllable at $\bar{x}_j$*, they are* controllable.

If the dynamics are not controllable at $\bar{x}_j$, we introduce the closed sets $\mathcal{X}_{j,-} = \{x \in \mathcal{X} : x \leq \bar{x}_j\}$ and $\mathcal{X}_{j,+} = \{x \in \mathcal{X} : x \geq \bar{x}_j\}$, and we decompose the optimisation problem into two coupled optimisation problems on $\mathcal{X}_{j,-}$ and $\mathcal{X}_{j,+}$ respectively.

Following Barles et al. (2014), we introduce the state-constrained value function V_-^{sc} and V_+^{sc} for, respectively, the optimisation problems where the state restriction $y(t) \in \mathcal{X}_{j,-}$, respectively $y(t) \in \mathcal{X}_{j,+}$, is required to hold for all t.

Additionally, we introduce $V_{\mathcal{J}}^{\text{sc}}(\bar{x}_j)$ for the state-constrained value function of the optimisation problem with the restriction $y(t) = \bar{x}_j$ for all t, where only action schedules that satisfy $\mathbf{f}(\bar{x}_j, a) = 0$ are admitted. Finally for $k \in \{-, \mathcal{J}, +\}$ we set $V_k^{\text{sc}} = -\infty$ if there is no trajectory starting at x that satisfies the particular state constraint.

The following result establishes necessary and sufficient conditions for the continuity of the value function interfaces. To formulate it, we introduce the notions of one-sided semi-repellers and semi-attractors.

Definition 16. *A point $\bar{x}_j \in \mathcal{J} \cup \partial X$ is a* left semi-repeller *if $f_j(\bar{x}_j, q) \leq 0$ for all $q \in \mathcal{Q}$, and a* right semi-repeller *if $f_{j+1}(\bar{x}_j, q) \geq 0$ for all $q \in \mathcal{Q}$.*

Likewise, $\bar{x}_j \in \mathcal{J} \cup \partial X$ is a left semi-attractor *if $f_j(\bar{x}_j, q) \geq 0$ for all $q \in \mathcal{Q}$, and a* right semi-attractor *if $f_{j+1}(\bar{x}_j, q) \leq 0$ for all $q \in \mathcal{Q}$.*

Proposition C.1. *Let $\bar{x}_j \in \mathcal{J} \cup \partial \mathcal{X}$ and V the value function. Then the following hold.*

(i) $V(\bar{x}_j) = \max\{V_-^{\text{sc}}(\bar{x}_j), V_{\mathcal{J}}^{\text{sc}}(\bar{x}_j), V_+^{\text{sc}}(\bar{x}_j)\}$.

(ii) If the dynamics are left (right) controllable at $\bar{x}_j$, then V is left (right) Lipschitz continuous in a left (right) neighbourhood of $\bar{x}_j$.

(iii) If $\bar{x}_j$ is a right semi-attractor, then V is right continuous at $\bar{x}_j$ and we have $V(\bar{x}_j) = \max\{V_-^{\text{sc}}(\bar{x}_j), V_{\mathcal{J}}^{\text{sc}}(\bar{x}_j)\}$.

(iv) If $\bar{x}_j$ is a left semi-attractor, then V is left continuous at $\bar{x}_j$ and we have $V(\bar{x}_j) = \max\{V_{\mathcal{J}}^{\text{sc}}(\bar{x}_j), V_+^{\text{sc}}(\bar{x}_j)\}$.

(v) If $\bar{x}_j$ is a left semi-repeller, then V is left continuous at $\bar{x}_j$ and $V(\bar{x}_j) = V_-^{\text{sc}}(\bar{x}_j)$.

(vi) The value function V is not continuous at $\bar{x}_j$ if and only if $\bar{x}_j$ is a right semi-repeller and $V_+^{\text{sc}}(\bar{x}_j) < \max\{V_-^{\text{sc}}(\bar{x}_j), V_{\mathcal{J}}^{\text{sc}}(\bar{x}_j)\}$.

The asymmetry in the result is a consequence of the fact that $u_x < 0$. The proof of this result is given in Appendix F.5.

The result motivates the following definition.

Definition 17. *Let $\mathscr{X}$ be a covering of $\mathcal{X}$, and let $\mathbf{f}$ be a dynamics defined on $\mathcal{X}$. The class $\mathscr{G}$ consists of functions $W : \mathcal{X} \to \mathbb{R}$ such that*

(i) W is decreasing;

(ii) W is left continuous everywhere;

(iii) W is continuous on $\mathring{\mathcal{X}}_j$ for all $j \in \{1, \dots, J\}$;

(iv) if W is not continuous at $\bar{x}_j$, then $j \in \{1, \dots, J-1\}$ and this point is a right semi-repeller under the dynamics $\mathbf{f}$.

The following result is central.

Theorem C.1. *The value function V is the unique viscosity solution of the Hamilton–Jacobi–Bellman equation* (16) *in the class $\mathscr{G}$.*

The theorem is a consequence of the following comparison result.

Proposition C.2. *Let $v \in \mathscr{G}$ and $w \in \mathscr{G}$ be respectively a supersolution and a subsolution of* (16). *Then $v(x) \geq w(x)$ for all $x \in \mathcal{X}$ where v and w are continuous.*

To prove this proposition, we need a number of technical results. The first gives the subsolution version of condition (18) for supersolutions. The subsolution version either holds, or an alternative property must be true. The result is adapted from Barles et al. (2013, Theorem 3.3): its proof is given in Section F.6.

Proposition C.3. *Let $w \in \mathscr{G}$ be a subsolution of* (16) *and $\bar{x} = \bar{x}_j \in \mathcal{I}$.*

Then either of the following two statements holds.

A. $\rho w(\bar{x}) - H^{\mathcal{I}}(\bar{x}) \leq 0$

B. (i) If w is continuous at $\bar{x}$, then there is a constant $\eta > 0$, an index $\ell \in \{j, j+1\}$, and a sequence $x_k \to \bar{x}$ such that $x_k \in \mathcal{X}_\ell$ for all k, $w(x_k) \to w(\bar{x})$ as $k \to \infty$, and for each k there is a trajectory–action pair (y_k, a_k) such that $y_k(0) = x_k$, $y_k(t) \in \mathcal{X}_\ell$ for all $t \in [0, \eta]$ and

$$w(x_k) \leq \int_0^\eta u(y_k(t), a_k(t)) \exp(-\rho t)\, \mathrm{d}t + w(y_k(\eta)) \exp(-\rho\eta)$$

(ii) If w is not continuous at $\bar{x}$, the previous statement holds with $\ell = j$.

The next result, proved in Section F.7, settles the continuous case of Proposition C.2.

Proposition C.4. *Let $v \in \mathscr{G}$ be a continuous viscosity supersolution, and $w \in \mathscr{G}$ a continuous viscosity subsolution of the HJB equation* (16)*. Then $v \geq w$ in $\mathcal{X}$.*

For the discontinuous case, we need the following technical result. Its proof is a standard viscosity test function argument, given in Section F.8.

Proposition C.5. *Let $j \in \{1, \dots, J-1\}$, $\mathcal{X}_{j,+} = \{x \in \mathcal{X} : x \geq \bar{x}_j\}$ and $v, w \in \mathscr{G}$ such that v and w are respectively a supersolution and a subsolution of* (16) *on the set $\mathcal{X}_{j,+} \backslash \{\bar{x}_j\}$, and either v or w is discontinuous at $\bar{x}_j$. Then v and w are also, respectively, a supersolution and a subsolution of* (16) *on $\mathcal{X}_{j,+}$.*

Proof of Proposition C.2. The proof proceeds by induction.

Let $x_{\min} \leq \hat{x}_1 < \hat{x}_2 < \dots < \hat{x}_{L-1} < x_{\max}$ be the right semi-repellers of the dynamics; then $\hat{x}_\ell = \bar{x}_{i_\ell}$ for $0 \leq i_0 < \dots < i_\ell < J$. Introduce for $\ell = 1, 2, \dots, L$ the intervals $\mathcal{X}^{(\ell)} = [\hat{x}_{L-\ell}, x_{\max}]$, as well as extensions v_ℓ, w_ℓ to $\mathcal{X}^{(\ell)}$ of the respective restrictions of v and w to $\mathcal{X}^{(\ell)} \backslash \{\hat{x}_{L-\ell}\}$ such that v_ℓ and w_ℓ are right continuous at $\hat{x}_{L-\ell}$: as $v, w \in \mathscr{G}$ these extensions are well-defined and unique. Proposition C.5 implies that v_ℓ and w_ℓ are respectively a supersolution and a subsolution of (16) on $\mathcal{X}^{(\ell)}$.

The induction hypothesis is that the inequality $v_\ell(x) \geq w_\ell(x)$ holds for all $x \in \mathcal{X}^{(\ell)}$ where v_ℓ and w_ℓ are continuous.

The functions v_1 and w_1 are continuous on $\mathcal{X}^{(1)}$. Proposition C.4 then implies the induction hypothesis for $\ell = 1$.

Assuming that the hypothesis is true for $\ell \geq 1$, consider $\Delta_{\ell+1} = w_{\ell+1} - v_{\ell+1}$. By the hypothesis, $\Delta_{\ell+1} \leq 0$ on $\mathcal{X}^{(\ell)} \backslash \{\hat{x}_{L-\ell}\}$. If $\Delta_{\ell+1} \leq 0$ on $\mathcal{X}^{(\ell+1)}$, there is nothing left to prove. If not, then $\Delta_{\ell+1}$ takes a positive maximum $M = \Delta_{\ell+1}(\bar{x}) > 0$ at a point $\bar{x} \in [\hat{x}_{L-(\ell+1)}, \hat{x}_{L-\ell}]$, as the interval is compact and $\Delta_{\ell+1}$ restricted to this interval is continuous.

If $\bar{x} \neq \hat{x}_{L-\ell}$, the same arguments used in the proof of Proposition C.4 can be used to derive a contradiction to the statement that $M > 0$. Hence we may assume that $\bar{x} = \hat{x}_{L-\ell}$. As this is an interface point, Proposition C.3, on which Proposition C.4 is based, applies.

If $v_{\ell+1}$ and $w_{\ell+1}$ are continuous at $\bar{x}$, the argument of Proposition C.4 for interface points again produces a contradiction.

If $v_{\ell+1}$ is discontinuous at $\bar{x}$, but $w_{\ell+1}$ is not, we have that $M = w_{\ell+1}(\bar{x}) - v_{\ell+1}(\bar{x}) > 0$ and, since $\Delta_{\ell+1} \leq 0$ if $x > \bar{x}$, also that $\lim_{x \downarrow \bar{x}}(w_{\ell+1}(x) - v_{\ell+1}(x)) \leq 0$, which implies that $v_{\ell+1}(\bar{x}) < \lim_{x \downarrow \bar{x}} v_{\ell+1}(\bar{x})$. But then $v_{\ell+1}$ cannot be an element of $\mathscr{G}$.

Finally, if $w_{\ell+1}$ is discontinuous at $\bar{x}$, the argument of Proposition C.4 for interface points holds again, as in Alternative B the sequence elements x_k satisfy $x_k \leq \bar{x}$ for all k.

We conclude that the induction hypothesis also holds for $\ell + 1$, and therefore for all $1 \leq \ell \leq L$. This completes the proof. □

Proof of Theorem C.1. By Theorem B.1, V is a viscosity solution to (16), and according to Proposition C.1, it is in $\mathscr{G}$.

Assume $W \in \mathscr{G}$ is another viscosity solution to (16). As V is a supersolution and W is a subsolution, by Proposition C.2 we have that $V \geq W$ at all points in $\mathcal{X}$ where V and W are continuous. Interchanging V and W yields also that $V \leq W$ at all points of continuity. As V and W are both left continuous everywhere, this shows that $V = W$ for all $x \in \mathcal{X}$. □

D Existence of the best response map

This section shows that the *best response map* $\phi_i = \mathscr{B}_i(\phi_{-i})$ is well-defined for all profiles ϕ_{-i} with adapted covering $\mathscr{X}$ that are in the complement of a set $\mathscr{E}_{\mathscr{X},-i}$, and it shows that the latter set is 'shy'—small in a topological as well as a measure-theoretical sense. The map $\mathscr{B}_i$ gives the Markovian best response of player i to the strategy profile ϕ_{-i} of the other players. That is, given the profile ϕ_{-i}, the strategy ϕ_i is the feedback strategy such that for every initial state the Markovian action schedule induced by ϕ_i maximises the total payoff U_i, given the dynamics $\mathbf{f}_i$.

The construction starts with a general profile $\phi_{-i} \in \mathscr{S}_{-i}$. In the first step the value function of player i is shown to be differentiable on a dense set of points. The second step improves the regularity: the value function has to be real analytic on non-constant optimal state orbits, and compact intervals not containing interface points intersect with only finitely many of these orbits. Restricted to such a compact interval, the set of points at which the value function is not real analytic is shown to be discrete, hence finite. The associated strategy can fail to be in $\mathscr{S}_i$ only if the points of non-analyticity accumulate on an interface point: the final step of the proof is to show that this only occurs for a shy set of profiles ϕ_{-i}.

D.1 Notations

In this section we work in a fixed interval $\mathcal{X}_j$. We therefore fix $j \in \{1, \ldots, J\}$ and drop this index for the sake of readability. Hence, in the whole section, unless announced differently, $f(x, q)$ stands for $f_j(x, q_j)$, which in turn stands for $f_{i,j}(x, q_{i,j})$, etc.

We introduce a number of auxiliary quantities. The functions $p_\ell, p_u : \mathcal{X} \to \mathbb{R}$ are given as

$$p_\ell(x) = -u_q(x, q_\ell)/f_q(x, q_\ell), \quad p_u(x) = -u_q(x, q_u)/f_q(x, q_u).$$

We have sets

$$\mathcal{P}_\ell = \{(x, p) \,:\, p \leq p_\ell(x)\}, \quad \mathcal{P}_u = \{(x, p) \,:\, p \geq p_u(x)\}, \quad \text{and} \quad \mathcal{P}_{\text{int}} = \mathcal{X} \times \mathbb{R} \backslash (\mathcal{P}_\ell \cup \mathcal{P}_u).$$

With these definitions, the maximiser $q = q^*(x, p)$ of $u(x, q) + pf(x, q)$ equals q_ℓ if $(x, p) \in \mathcal{P}_\ell$, q_u if $(x, p) \in \mathcal{P}_u$, and it takes a value in $\mathring{\mathcal{Q}}$ if $(x, p) \in \mathcal{P}_{\text{int}}$.

The boundaries of these sets are the *switching manifolds*

$$\mathcal{S}_\ell = \{(x, p) \,:\, p = p_\ell(x)\}, \quad \mathcal{S}_u = \{(x, p) \,:\, p = p_u(x)\}, \quad \mathcal{S} = \mathcal{S}_\ell \cup \mathcal{S}_u.$$

D.2 Differentiability of the value function in a dense set

We partition the interior of $\mathcal{X}$ into three sets $\mathcal{X}_L(0)$, $\mathcal{X}_M$ and $\mathcal{X}_N$, based on whether $H(x, p)$ has a single minimiser with respect to p or not, and, if not, whether $H(x, p) = \rho V(x)$ is locally solvable for p or not.

Definition 18. *Let $c \geq 0$ and $k = 0, 1, 2 \ldots.$ Introduce the sets*

(i) $\mathcal{D}_x = \{x \,:\, V$ is k times differentiable at $x\}$;

(ii) $\mathcal{X}_L(c) = \{x : H_p(x,p) < -c \text{ if } p \leq p_\ell(x) \text{ and } H_p(x,p) > c \text{ if } p \geq p_u(x)\}$;

(iii) $\mathcal{X}_M = \{x : \exists!\, p \text{ such that } \rho V(x) = H(x,p) \text{ and } H_p(x,p) \neq 0\}$;

(iv) $\mathcal{X}_N = \{x : \exists\, p \text{ s.t. } p \leq p_\ell(x) \text{ or } p \geq p_u(x), \rho V(x) = H(x,p), \text{ and } H_p(x,p) = 0\}$.

The sets $\mathcal{X}_L(0)$, $\mathcal{X}_M$ and $\mathcal{X}_N$ are mutually disjoint and satisfy $\mathcal{X}_L(0) \cup \mathcal{X}_M \cup \mathcal{X}_n = \mathring{\mathcal{X}}$.

Proposition D.1. *The value function V is differentiable almost everywhere on $\mathcal{X}_L(0)$, and it is real analytic on $\mathcal{X}_M \cup \mathring{\mathcal{X}}_N$.*

Proof. Take $c > 0$. As V is continuous in $\mathring{\mathcal{X}}$ and a supersolution of the Hamilton–Jacobi–Bellman equation, for $x \in \mathcal{X}_L(c)$ and $p \in D^- V(x)$, we have for every $q \in \mathcal{Q}$ that

$$\rho V(x) \geq H(x,p) \geq u(x,q) + pf(x,q),$$

which implies, since $|V(x)| \leq \|u\|_\infty/\rho$ for all x, that

$$pf(x,q) \leq \rho V(x) - u(x,q) \leq 2\|u\|_\infty.$$

Taking $q = q_u$, and using that $f(x,q_u) = H_p(x,p_u(x)) > c$, we find $p < 2\|u\|_\infty/c$. Similarly, for $q = q_\ell$, we have $f(x,q_\ell) = H_p(x,p_\ell(x)) < -c$ and $p > -2\|u\|_\infty/c$. We conclude that if $p \in D^- V(x)$, then $|p| < 2\|u\|_\infty/c$.

Bardi and Capuzzo-Dolcetta (2008, Remark II.5.16) now implies that V is Lipschitz continuous on $\mathcal{X}_L(c)$ with Lipschitz constant $2\|u\|_\infty/c$; Rademacher's theorem (Clarke et al., 1998, Chapter 3, Corollary 4.19) subsequently ensures almost everywhere differentiability of V on $\mathcal{X}_L(c)$, and hence on $\mathcal{X}_L(0) = \bigcup_{c>0} \mathcal{X}_L(c)$, proving the first part of the statement.

Next, consider $x \in \mathcal{X}_M$. By the implicit function theorem the solution $p = \kappa(x,w)$ of $w = H(x,p)$ is locally real analytic.

Assume that $H_p(x,\kappa(x,w)) > 0$, the other situation being similar. As $x \in \mathcal{X}_M$, this implies that $H_p(x,p) \geq 0$ for all p. If $p \in D^+ V(x)$, the subsolution property of V implies

$$H(x,\kappa(x,\rho V(x))) = \rho V(x) \leq H(x,p)$$

and therefore $p \geq \kappa(x,\rho V(x))$, by convexity of $H(x,p)$ in p; similarly, the supersolution property implies for $p \in D^- V(x)$ that $p \leq \kappa(x,\rho V(x))$. Again using Bardi and Capuzzo-Dolcetta (2008, Remark II.5.16), it follows that $V(x)$ is a classical solution of $V'(x) =$

$\kappa(x, \rho V(x))$. Since κ is real analytic, it follows that V is a real analytic solution of $\rho V = H(x, V')$ in $\mathcal{X}_M$.

Finally, for $x \in \mathring{\mathcal{X}}_N$, $\rho V(x) = u(x, q_\ell)$ or $\rho V(x) = u(x, q_u)$, and V is again real analytic. This proves the second part. $\square$

As a corollary of this result, we obtain

Proposition D.2. *The set $\mathcal{D}_1$ is dense in $\mathring{\mathcal{X}}$.*

Proof. If this were not the case, there is a point $\bar{x} \in \partial\mathcal{X}_N$ such that V is not differentiable for any point in an open interval I with positive length containing $\bar{x}$. As $\bar{x}$ is a boundary point of $\mathcal{X}_N$, there is a point $\tilde{x} \in I \backslash \mathcal{X}_N$. Hence the intersection $I \cap (\mathcal{X}_L(0) \cup \mathcal{X}_M)$ is nonempty. But this intersection is open, and therefore it contains a positive measure subset of points where V is differentiable, which is a contradiction. $\square$

D.3 Canonical trajectories

To extend the domain of differentiability of the value function, we show that differentiability is carried forward along optimal orbits by the costate dynamics. This result is closely related to the Pontryagin Maximum Principle in the finite horizon context, the difference being that we here have initial values rather than terminal values for the costate equation. The result, whose proof is given in Section F.9, is an adaptation of Cannarsa and Frankowska (1991, Theorem 3.3) to the present context.

Proposition D.3. *Let (y^*, a^*) be an optimal trajectory–action pair with initial point $x \in \mathcal{D}_1$ and let $T \geq 0$ be such that $y^*(t) \in \mathring{\mathcal{X}}$ for all $0 \leq t \leq T$.*

Let moreover p^* be the solution of

$$\dot{p}(t) = \rho p(t) - u_x(y^*(t), a^*(t)) - p(t) f_x(y^*(t), a^*(t)), \qquad p(0) = V'(x). \tag{20}$$

Then for every $0 \leq t \leq T$ we have that $y^*(t) \in \mathcal{D}_1$, $V'(y^*(t)) = p^*(t)$, and $a^*(t) = q^*(y^*(t), p^*(t))$.

Proposition D.3 can be expressed in the more familiar form that an optimal trajectory necessarily satisfies the *canonical equations*

$$\dot{y}(t) = H_p(y(t), p(t)), \qquad \dot{p}(t) = \rho p(t) - H_x(y(t), p(t)),$$

with $y(0) = x$, $p(0) = V'(x)$. This motivates the following definition:

Definition 19. *The* canonical vector field $X : \mathcal{X} \times \mathbb{R} \to \mathbb{R}^2$ *is given as* $X(x, p) = (H_p(x, p), \rho p - H_x(x, p))$. *A* canonical trajectory *is a trajectory of* X. *An* optimal canonical trajectory *is a canonical trajectory* (y, p) *such that* y *is an optimal trajectory.*

The next proposition, proved in Section F.10, states properties of non-constant optimal trajectories: only at their initial states the value function may be non-differentiable; they are monotone; and they correspond to optimal canonical trajectories, which moreover converge to steady states of X in $\mathcal{X}$, or leave the interval $\mathcal{X}$ $(= \mathcal{X}_j)$.

Proposition D.4. *Let* (y^*, a^*) *be a non-constant optimal trajectory–action pair.*

(i) $y^*(t) \in \mathcal{D}_1$ *for all* $t > 0$ *such that* $y^*(t) \in \mathring{\mathcal{X}}$*;*

(ii) there is a unique initial point (x, p_0) *such that the canonical trajectory* (y, p) *through this point satisfies* $y^*(t) = y(t)$ *for all* $t \geq 0$*;*

(iii) either $\dot{y}^*(t) > 0$ *for all* $t \geq 0$*, or* $\dot{y}^*(t) < 0$ *for all* $t \geq 0$*;*

(iv) either $T = \inf\{t > 0 : y^*(t) \notin \mathcal{X}\}$ *is finite, or there is a steady state* $(\bar{x}, \bar{p})$ *such that* $(y^*(t), p^*(t)) \to (\bar{x}, \bar{p})$ *as* $t \to \infty$*.*

D.4 Markovian best responses

This section constructs the piecewise real analytic Markovian best response strategy. For the remainder of this section, we reinstate the full indexed notation.

Definition 20. *An* optimal orbit *is an interval* $I \subset \mathcal{X}$ *such that there is* $\mathcal{T} = [0, \infty)$ *or* $\mathcal{T} = \mathbb{R}$ *and a state trajectory* $y : \mathcal{T} \to I$ *with the property that for every* $x \in I$ *there is* $\tau \in \mathcal{T}$ *such that* $y_\tau(t) = y(\tau + t)$ *satisfies* $y_\tau(0) = x$ *and* y_τ *is an optimal trajectory. If* I *consists of a single point, it is an* optimal steady state*; if* I *has positive length, it is an* optimal non-constant orbit.

The next two results, proved in Section F.11, give the structure of the set of optimal non-constant orbits in $\mathcal{X}_j$: there are at most countably many, and they can only accumulate on the end points of $\mathcal{X}_j$. Moreover, restricted to the interior of a non-constant orbit, a best response exists and is real analytic.

Proposition D.5. *A compact interval $C \subset \mathring{\mathcal{X}}_j$ intersects finitely many optimal non-constant orbits.*

Definition 21. *Let (y,p) be a canonical trajectory. A* switching point *of this trajectory is a point $(y(t_0), p(t_0)) \in \mathcal{S}_b$, $b \in \{\ell, u\}$, such that in every neighbourhood of t_0 there are t_1, t_2 with $(y(t_1), p(t_1)) \in \mathcal{P}_{int}$ and $(y(t_2), p(t_2)) \in \mathcal{P}_b$. Switching points for non-constant optimal state trajectories and optimal orbits are defined as switching points of their associated canonical trajectories.*

For the remainder of the section, we introduce for an interval $I \subset \mathcal{X}$ the notation $\mathscr{S}_i(I)$ for the strategy space $\mathscr{S}_i$ with the interval $\mathcal{X}$ replaced by I.

Proposition D.6. *If the open interval $I \subset \mathcal{X}$ is part of a non-constant optimal orbit, does not contain switching points and satisfies $\phi_{i,j}(x) = q_i^*(x, V_i'(x))$ for all $x \in I$, then V_i is real analytic on $J = I \cap \mathring{\mathcal{X}}_j$ and $\phi_{i,j} \in \mathscr{S}_i(J)$ for all j.*

By piecing together the results for constant and non-constant optimal orbits, we obtain the existence of Markovian best responses.

Proposition D.7. *Let $C \subset \mathring{\mathcal{X}}_j$ be a compact interval. Then V is piecewise real analytic on C, the function defined as $\phi_{i,j}(x) = q_i^*(x, V_i'(x))$ for all $x \in \mathcal{D}_1 \cap C$ can be extended to a strategy $\phi_{i,j} \in \mathscr{S}_i(C)$, and $\phi_{i,j}$ is the unique Markovian best response to $\phi_{-i,j}$ on C. Each point of non-analyticity of $\phi_{i,j}$ is contained in the closure of some non-constant orbit.*

Proof. By Proposition D.5, the interval C intersects only finitely many non-constant optimal orbits $I_1, \ldots, I_m$. The complement of the intersection consists of finitely many compact intervals $N_1, \ldots, N_n$ containing only optimal steady states.

If $x \in \mathring{N}_k$, then $V_i(x) = u_i(x, \bar{q}_i(x))/\rho$, where $q = \bar{q}_i(x)$ solves $f_{i,j}(x,q) = 0$. By the implicit function theorem, the function $\bar{q}_i$, and hence V_i and $\phi_{i,j}$, are real analytic in this set, showing the last statement of the proposition. Proposition D.6 implies that V_i is real analytic on $\mathring{I}_k \cap \mathring{\mathcal{X}}_j$ and $\phi_{i,j} \in \mathscr{S}_i(\mathring{I}_k \cap \mathring{\mathcal{X}}_j)$ for each k. We therefore conclude that V_i is piecewise real analytic on C and $\phi_{i,j} \in \mathscr{S}_i(C)$.

It is clear that $\phi_{i,j}$ is uniquely determined: it remains to show that it is a best response. Let V_i be differentiable at a point $x \in \mathcal{X}_j$, and let (y^*, a^*) be the optimal trajectory–action pair with initial point x. Then Proposition D.3 implies that

$$a^*_{i,j}(t) = q^*_{i,j}(y^*(t), V_i'(y^*(t)))$$

for all $t \geq 0$. This implies the compatibility condition (9). $\square$

Proof of Theorem 1. Take $\phi_{-i} \in \mathscr{S}_{-i}$, and let $\mathscr{X}$ be a covering adapted to ϕ_{-i}. For the duration of this proof, we write $\mathscr{F} = \mathscr{F}^{N-1}_{\mathscr{X}}$. A profile $\phi_{-i} \in \mathscr{S}_{\mathscr{X},-i} = \mathscr{F} \cap \mathscr{S}_{-i}$ then satisfies $\phi_{-i,j} \in \mathscr{S}_{-i}(\mathcal{X}_j)$. Let $\phi_{i,j} : \mathring{\mathcal{X}}_j \to \mathbb{R}$ be the unique Markovian best response to ϕ_{-i} on compact subsets of $\mathring{\mathcal{X}}_j$, given by Proposition D.7.

If $\phi_{i,j}$ has infinitely many points of non-analyticity, by Propositions D.6 and D.7 they have to be either endpoints of non-constant optimal orbits or switching points. Every compact subinterval of $\mathring{\mathcal{X}}_j$ contains only finitely many of these: hence points of non-analyticity have to accumulate on a boundary point of $\mathcal{X}_j$. In particular, at such a boundary point the canonical vector field either is tangent to $\mathcal{P}_\ell$ or $\mathcal{P}_u$ or vanishes. We shall show that the set $\mathscr{E} \subset \mathscr{S}_{-i}$ of profiles ϕ_{-i} such that $\phi_{i,j}$ has this latter property for some j is shy.

The space $\mathscr{F}$ is a complete metric linear space. Let $w^{(1)}(x) = 1$ and $w^{(2)}(x) = \prod_{j=0}^{J}(x - \bar{x}_j)$ and introduce $\psi^{(k)} \in \mathscr{F}$ as $\psi^{(k)}(x) = (w^{(k)}(x), \ldots, w^{(k)}(x))$, for $k = 1, 2$. Denote by $\mathscr{L}$ the two-dimensional subspace of $\mathscr{F}$ spanned by the $\psi^{(k)}$, and let meas be the Lebesgue measure on $\mathscr{L}$. Let $X_{\phi_{-i,j}}$ be the canonical vector field on $\mathcal{X}_j \times \mathbb{R}$ associated to $\phi_{-i,j}$.

Let $S_0 = \left\{\phi_{-i} \in \mathscr{F} \,:\, X_{\phi_{-i}} = 0 \text{ at a boundary point of } \mathcal{X}_j\right\}$. For $b \in \{\ell, u\}$ and $n_b(x)$ the normal vector to $\mathcal{S}_b$ at $(x, p_b(x))$ pointing out of $\mathcal{P}_{\mathrm{int}}$, let

$$SW_b = \left\{\phi_{-i} \in \mathscr{F} \,:\, n_b \cdot X_{\phi_{-i}} = 0 \text{ at a boundary point of } \mathcal{X}_j\right\}.$$

Note that $\mathscr{E} \subset S_0 \cup SW_\ell \cup SW_u \equiv \mathscr{N}$.

We prove shyness for SW_b, the proof for S_0 being similar. Take $\phi_{-i} \in \mathscr{S}_{\mathscr{X},-i}$ and $\hat{\phi}_{-i} \in SW_b$. The point $\hat{\phi}_{-i} - \phi_{-i}$ is in $\mathscr{L}$ if $\hat{\phi}_{-i} = \phi_{-i} + \lambda_1\psi^{(1)} + \lambda_2\psi^{(2)}$ for some

$\lambda_1, \lambda_2 \in \mathbb{R}$. We have $\hat{\phi}_{-i} \in SW_b$ if

$$0 = n_b \cdot X_{\hat{\phi}_{-i}} = n_{b,1} f(x, q_i^*, \hat{\phi}_{-i}) + n_{b,2}\Big(\rho p_{i,b}(x) - (u_i)_x(x, q_i^*)\Big) \\ - n_{b,2} p_{i,b}(x)\Big(f_x(x, q_i^*, \hat{\phi}_{-i}) + f_{q_{-i}}(x, q_i^*, \hat{\phi}_{-i})\hat{\phi}'_{-i}\Big). \tag{21}$$

For $\bar{x} \in \partial \mathcal{X}_j$ we have $(\psi^{(1)})'(\bar{x}) = 0$ and $\psi^{(2)}(\bar{x}) = 0$. Equation (21) reduces to

$$0 = n_b \cdot X_{\phi_{-i}+\lambda_1\psi^{(1)}} - \lambda_2 n_{b,2}(\bar{x}) p_{i,b}(\bar{x}) w'(\bar{x}) \sum_{k\neq i} \frac{\partial f}{\partial q_k}(\bar{x}, q^*(\bar{x}, p_i, b(\bar{x})), \phi_{-i}(\bar{x})). \tag{22}$$

Consider first the situation that $p_{i,b}(\bar{x}) \neq 0$. Since $n_{b,2}(\bar{x})$ and $w'(\bar{x})$ are both nonzero, and $\frac{\partial f}{\partial q_k} > 0$ for all k by Assumption 1, the solutions (λ_1, λ_2) of equation (22) are located on a graph $\lambda_2 = \lambda_2(\lambda_1)$. There are $2J$ such graphs, one for every endpoint of every $\mathcal{X}_j$. Hence the set $SW_b \cap \phi_{-i} + \mathscr{L}$ has measure zero in this case.

If $p_{i,b}(\bar{x}) = 0$, equation (22) reads as $n_{b,1} f(x, q_i^*, \phi_{-i} + \lambda_1 w^{(1)}) - n_{b,2}(u_i)_x(x, q_i^*) = 0$. If $n_{b,1} = 0$, this has no solution by Assumption 2; if $n_{b,1} \neq 0$, by Assumption 1, this equation has a unique constant solution $\lambda_1 = \lambda_1(\lambda_2)$. As before, we infer that $SW_b \cap \phi_{-i} + \mathscr{L}$ has measure zero, and we conclude that SW_b is shy.

Unions and subsets of shy sets are shy (Hunt et al., 1992). Hence $\mathscr{N}$ and $\mathscr{E} \subset \mathscr{N}$ are shy. The complement of a shy set is dense. As $\mathscr{N}$ is closed, it follows that its complement is also open, and that $\mathscr{N}$, and therefore $\mathscr{E}$, are nowhere dense.

Take $\phi_{-i} \in \mathscr{S}_{\mathscr{X},-i} \backslash \mathscr{N}$. Then the strategy profile $\phi_{-i,j}$ has only finitely many points of non-analyticity on $\mathring{\mathcal{X}}_j$. It remains to show that the limit of $\phi_{i,j}(x)$ and $\phi'_{i,j}(x)$ exist as x tends to the boundary of $\mathcal{X}_j$.

Let $I \subset \mathcal{X}_j$ be a maximal open interval on which $\phi_{i,j}$ is real analytic and that contains a boundary point $\bar{x}$ of $\mathcal{X}_j$. Set $\bar{p} = \lim_{x\to\bar{x}, x\in I} V'(x)$. If $(x, V'(x))$ is contained in $\mathcal{P}_b$ for all $x \in I$, $b = \ell, u$, the best response $\phi_{i,j}(x)$ equals q_b and can be extended to a continuously differentiable function on the closure of I.

If $(x, V'(x)) \in \mathcal{P}_{\text{int}}$, and if $\bar{z} = (\bar{x}, \bar{p})$ is not a steady state, we have, since $(X_{\phi_{-i,j}})_1(\bar{x}) \neq 0$, that there is a continuously differentiable extension of the canonical trajectory through $\bar{z}$ to a neighbourhood of $\bar{z}$, giving rise to a continuously differentiable extension of $\phi_{i,j}$ to $\bar{x}$. If $\bar{z}$ is a steady state, the graph $(x, V'(x))$ is tangent to a, necessarily one-dimensional, eigenspace of $DX_{\phi_{-i,j}}(\bar{z})$ at $\bar{z}$ and can be extended as a continuously differentiable function as well. This shows that $\phi_i \in \mathscr{S}_i$. $\square$

Corollary 1. *Let $\mathscr{P} \subset \mathscr{S}_{-i}$ be the set of profiles ϕ_{-i} such that for every j the canonical vector field $X_{\phi_{-i,j}}$ of player i has only finitely many steady states and finitely many switching points. Then for every $\phi_{-i} \in \mathscr{P}$ there is a unique best response $\phi_i \in \mathscr{S}_i$.*

E Characterisation of the best response map

We continue not to indicate the player index i. This section provides the proof of Theorem 2. Before giving this proof, we need to collect information about V^ϕ.

Introduce for $x \in \mathfrak{X}_j$ the functions

$$f_j^\phi(x) = f_j(x, \phi_j(x)), \quad u_j^\phi(x) = u_i(x, \phi_j(x)), \quad H_j^\phi(x,p) = u_j^\phi(x) + p f_j^\phi(x).$$

In terms of these functions, set

$$\mathbf{u}^\phi(x) = \begin{cases} \sum_{j=1}^J \mathbf{1}_{\mathfrak{X}_j}(x) u_j^\phi(x), & x \in \mathfrak{X}\backslash\mathfrak{I}, \\ \mu_j^\phi(x) u_j^\phi(x) + (1-\mu_j^\phi(x)) u_{j+1}^\phi(x), & x = \bar{x}_j \in \mathfrak{I}, \end{cases}$$

and

$$\mathbf{f}^\phi(x) = \begin{cases} \sum_{j=1}^J \mathbf{1}_{\mathfrak{X}_j}(x) f_j^\phi(x), & x \in \mathfrak{X}\backslash\mathfrak{I}, \\ \mu_j^\phi(x) f_j^\phi(x) + (1-\mu_j^\phi(x)) f_{j+1}^\phi(x), & x = \bar{x}_j \in \mathfrak{I}. \end{cases}$$

Then the value V^ϕ accruing to player i under the profile ϕ has the following properties, proved in Appendix F.12.

Proposition E.1. *Assume $\phi \in \mathscr{S}$ and $i \in \{1, \dots, N\}$. Then there is a covering $\mathscr{X} = \{\mathfrak{X}_j\}_{j=1}^J$ of X, with $\mathfrak{X}_j = [x_{j-1}, x_j]$, adapted to ϕ, and a finite set $\mathcal{E} \subset \mathfrak{X}$, such that the following hold.*

(i) $\|V^\phi\|_\infty \leq \max\{\|u\|_\infty/\rho, \|\beta\|_\infty\}$.

(ii) *If $x = \bar{x}_0$, then either $V^\phi(x) = \beta(x)$ or $f_1^\phi(x) \geq 0$.*
Likewise, if $x = \bar{x}_J$, then either $V^\phi(x) = \beta(x)$ or $f_J^\phi(x) \leq 0$.

(iii) *If $x = \bar{x}_j \in \mathcal{I}$, V^ϕ is continuous at x, and either $f_j^\phi(x) < 0 < f_{j+1}^\phi(x)$ or $f_j^\phi(x) > 0 > f_{j+1}^\phi(x)$, then $\rho V^\phi(x) = \mathbf{u}^\phi(x)$.*

(iv) *The function V^ϕ is continuous in $\mathcal{X} \backslash \mathcal{I}$ and real analytic in $\mathcal{X} \backslash (\mathcal{I} \cup \mathcal{E})$.*

(v) *If $x \in \mathcal{X}_j \backslash (\mathcal{I} \cup \mathcal{E}_j)$, then $\rho V^\phi(x) = H_j^\phi(x, (V^\phi)'(x))$, while if $x \in \mathcal{E}_j$, then $f_j^\phi(x) = 0$ and $\rho V^\phi(x) = u_j^\phi(x)/\rho$.*

(vi) *If $x \in \mathcal{X}$ and $z \uparrow x$ ($z \downarrow x$), there is a value $p_-, (p_+) \in \mathbb{R} \cup \{-\infty, \infty\}$, such that $(V^\phi)'(z) \to p_-$ ($(V^\phi)'(z) \to p_+$).*

Proof of Theorem 2. We have to show that Conditions (i)–(vi) of the theorem imply that V^ϕ is in class $\mathscr{G}$ and that it is a viscosity solution of (16). Theorem C.1 then implies that $V^\phi = V$, and hence that $\phi = \phi_i$ is a best response to ϕ_{-i}.

Conversely, we have to show that if ϕ is a best response, then Conditions (i)–(vi) hold true.

E.1 Notations

We recall the notation $q_j^*(x, p)$ for the maximiser of $q \mapsto u(x, q) + pf_j(x, q)$ over $\mathcal{Q}$, the local Hamilton functions $H_j(x, p) = u(x, q_j^*(x, p)) + pf_j(x, q_j^*(x, p))$, as well as $p_{j,b}(x) = -u_q(x, q_b)/(f_j)_q(x, q_b)$ for $b \in \{\ell, u\}$.

For a given x, we write the left and right limits of $(V^\phi)'$ at x as

$$p_- = \lim_{z \uparrow x} (V^\phi)'(z) \quad \text{and} \quad p_+ = \lim_{z \downarrow x} (V^\phi)'(z).$$

Proposition E.1(vi) ensures that these limits exist everywhere in $\mathcal{X}$, if we allow the possibility that the limits take the values $-\infty$ or ∞.

For $x \in \mathcal{I}_j$, introduce the further abbreviations $H_-(p) = H_j(x, p)$ and $H_+(p) = H_{j+1}(x, p)$.

E.2 Sufficiency

Assume that Conditions (i)–(vi) hold. Proposition E.1(iv), as well as Conditions (ii) and (iv) imply that $V^\phi \in \mathscr{G}$. We have to show that it is a viscosity solution of (16) for every $x \in \mathcal{X}$.

Subdifferentials and superdifferentials. For any point $x \in \mathcal{X}$ where V^ϕ is continuous, if $p_- < p_+$, then $D^-V^\phi(x) = [p_-, p_+]$ and $D^+V^\phi(x) = \emptyset$; similarly, if $p_+ < p_-$, then $D^-V^\phi(x) = \emptyset$ and $D^+V^\phi(x) = [p_+, p_-]$; finally if $p_- = p_+ = p$, then $D^-V^\phi(x) = D^+V^\phi(x) = \{p\}$. The final situation occurs if and only if V^ϕ is differentiable at x.

For a point $x \in \mathcal{X}$ at which V^ϕ is not continuous, Condition (ii) implies that

$$D^-(V^\phi)_*(x) = (-\infty, p_+] \quad \text{and} \quad D^+(V^\phi)^*(x) = (-\infty, p_-].$$

Interior of $\mathcal{X}_j$. Take first $x \in \mathring{\mathcal{X}}_j$ for $j \in \{1, \dots, J\}$: by Condition (iv) the function V^ϕ is continuous at x.

If V^ϕ is differentiable at x, set $p = (V^\phi)'(x)$: then Condition (i) and Proposition E.1(v) imply

$$H_j(x,p) = u(x,\phi(x)) + pf_j(x,\phi(x)) = H_j^\phi(x,p) = \rho V^\phi(x). \tag{23}$$

If V^ϕ is not differentiable at x, then $p_- \neq p_+$. Since ϕ is continuous at x, we have that $\phi(x) = q_j^*(x,p_-) = q_j^*(x,p_+)$, and therefore either $p_-, p_+ \le p_{j,\ell}(x)$ or $p_-, p_+ \ge p_{j,u}(x)$. Take $p \in D^-V^\phi(x) \cup D^+V^\phi(x)$: one of the two sets is empty. Then $q_j^*(x,p) = \phi(x) = q_b$, with $b \in \{\ell, u\}$. By Proposition E.1(v) we have $f_j(x,q_b) = 0$ and $\rho V^\phi(x) = u(x,q_b)$.

It follows that

$$\rho V^\phi(x) = u(x,q_b) + pf_j(x,q_b) = u(x,q_j^*(x,p)) + pf_j(x,q_j^*(x,p)) = H_j(x,p). \tag{24}$$

Equations (23) and (24) together show that (17) and (19) hold for all $x \in \mathring{\mathcal{X}}_j$.

Interface points at which V^ϕ *is continuous.* Take $x = \bar{x}_j \in \mathcal{I}$ with V^ϕ is continuous at x. If $D^-(V^\phi)(x)$ is nonempty, we have to show that

$$\rho V^\phi(x) \ge \min\{H_-(p), H_+(p)\}$$

for all $p \in D^-(V^\phi)(x)$. By continuity, $\rho V^\phi(x) = H_-(p_-) = H_+(p_+)$. Assume there is a point $\hat{p} \in (p_-, p_+)$ such that

$$H_-(p_-) = \rho V^\phi(x) < H_-(\hat{p}) \quad \text{and} \quad H_+(p_+) = \rho V^\phi(x) < H_+(\hat{p}).$$

By convexity of H_- and H_+, it follows that $\hat{f}_- := (H_-)_p(\hat{p}) > 0$ and $\hat{f}_+ := (H_+)_p(\hat{p}) < 0$. Hence there are $\lambda_-, \lambda_+ > 0$ such that $\lambda_- + \lambda_+ = 1$ and $\lambda_-\hat{f}_- + \lambda_+\hat{f}_+ = 0$. Set $u_- = u_j(x, q_j^*(x, \hat{p}))$ and $u_+ = u_{j+1}(x, q_{j+1}^*(x, \hat{p}))$. Condition (v) then implies that

$$\begin{aligned} \rho V^\phi(x) &\geq \lambda_- u_- + \lambda_+ u_+ = \lambda_- u_- + \lambda_+ u_+ + \hat{p}(\lambda_-\hat{f}_- + \lambda_+\hat{f}_+) \\ &= \lambda_- H_-(\hat{p}) + \lambda_+ H_+(\hat{p}) \geq \min\{H_-(\hat{p}), H_+(\hat{p})\} > \rho V^\phi(x), \end{aligned}$$

a contradiction, which proves (17) in this situation.

Next, assume that $D^+V^\phi(x)$ is nonempty. We have to show that

$$\rho V^\phi(x) \leq \max\{H_-(p), H_+(p)\}$$

for all $p \in D^+(V^\phi)(x) = [p_+, p_-]$. Assume, as before, that the relation does not hold for some $\hat{p} \in (p_+, p_-)$, that is

$$H_-(p_-) = \rho V^\phi(x) > H_-(\hat{p}) \quad \text{and} \quad H_+(p_+) = \rho V^\phi(x) > H_+(\hat{p}).$$

Convexity now implies that $f_- := (H_-)_p(p_-) > 0$ and $f_+ := (H_+)_p(p_+) < 0$. Since $f_- = f_j(x, \phi_j(x))$ and $f_+ = f_{j+1}(x, \phi_{j+1}(x))$, Condition (vi) implies that V^ϕ is differentiable at x, and we have therefore that $p_- = p_+$ and $\rho V^\phi(x) = H_-(p) = H_+(p)$ for all $p \in D^+V^\phi(x)$. Hence (19) holds for this case.

Interface points at which V^ϕ is not continuous. The next situation to consider is $x = \bar{x}_j \in \mathfrak{I}$ such that V^ϕ is not continuous at x.

To show that (17) holds in this case, assume that there is $\hat{p} \in D^-(V^\phi)_* = (-\infty, p_+]$ for which $\rho(V^\phi)_*(x) < \min\{H_-(\hat{p}), H_+(\hat{p})\}$. Then we have in particular that $\hat{p} < p_+$ and $H_+(p_+) = \rho(V^\phi)_*(x) < H_+(\hat{p})$. Convexity of H_+ implies that $(H_+)_p(\hat{p}) < 0$. This however contradicts Condition (iv), which implies that $(H_+)_p(p) \geq 0$ for all p.

Turning to (19), assume there is $\hat{p} \in D^+(V^\phi)^* = (-\infty, p_-]$ for which

$$\rho(V^\phi)^*(x) > \max\{H_-(\hat{p}), H_+(\hat{p})\}.$$

Then $\hat{p} < p_-$ and $H_-(p_-) = \rho(V^\phi)^*(x) > H_-(\hat{p})$. Invoking convexity of H_-, we obtain that $(H_-)_p(p_-) = f_j(x, \phi_j(x)) > 0$. Again, this is incompatible with Condition (iv).

Boundary points. We only consider the situation that $x = \bar{x}_0$, the other being entirely analogous. At x, we have that $D^+V^\phi(x) = [p_+, \infty)$ and $D^-V^\phi(x) = (-\infty, p_+]$. It follows from Condition (iv) that V^ϕ is continuous at x.

To prove (17) at x, assume that $V^\phi(x) < \beta(x)$. Condition (iii) then implies that $f_1(x, q) \geq 0$ for all $q \in \mathcal{Q}$. In particular $f(x, q^*(x, p)) = (H_+)_p(p) \geq 0$ for all p and $H_+(p)$ is non-decreasing in p. Since $\rho V^\phi(x) - H_+(p_+) = 0$ by continuity, it follows that $\rho V^\phi(x) - H_+(p) \geq 0$ for all $p \in (-\infty, p_+] = D^-V^\phi(x)$, which implies (17).

To show (19) at x, assume that $V^\phi(x) > \beta(x)$. By Proposition E.1(ii) and Condition (i), we have $f^\phi(x) = f(x, q^*(x, p_+)) = (H_+)_p(p_+) \geq 0$, and, by convexity of $H_+(p)$, it follows that $H_+(p) \geq H_+(p_+)$ for all $p > p_+$, implying that $\rho V(x) - H_+(p) \leq \rho V(x) - H_+(p_+) = 0$ for all $p \in D^+V^\phi(x)$.

Finally, Condition (v) implies (18). This concludes the proof of the sufficiency part.

E.3 Necessity

To prove the necessity of Conditions (i)–(vi) — reinstating the player index i for a moment — assume that ϕ_i is the best response to ϕ_{-i}. Then $V_i^\phi = V_i$ is the viscosity solution of (16).

Maximum principle. If $x \in \mathring{\mathcal{X}}_j$ and V is differentiable at x, then $D^-V(x) = D^+V(x) = \{V'(x)\}$, and (17) and (19) imply that $\rho V(x) = H_j(x, V'(x))$. Moreover, since $V = V^\phi$, we also have that $V'(x) = (V^\phi)'(x) =: p$ and $H_j(x, p) = H_j^\phi(x, p)$, which is equivalent to

$$u(x, \phi_j(x)) + pf(x, \phi_j(x), \phi_{-j}(x)) = \max_q \left(u(x, q) + pf(x, q, \phi_{-j}(x))\right),$$

and therefore implies Condition (i).

Monotonicity. Condition (ii) follows from Proposition A.3.

Boundary values. We show Condition (iii) for $x = \bar{x}_0$, the other case being analogous.

If $\rho V(x) < \rho\beta(x) = H_-(p)$, by (17) we have that $\rho V(x) - H_+(p) \geq 0$ for all $p \in D^- V(x) = (-\infty, p_+]$. By convexity of H_+, this implies that $(H_+)_p = f_+(x, q^*(x,p)) \geq 0$ for all $p \leq p_+$, which implies in particular that $f_+(x, q_\ell) \geq 0$, and hence $f_+(x,q) \geq 0$ for all $q \in \mathcal{Q}$.

Value discontinuities. To show Condition (iv), note that, since $V \in \mathscr{G}$, if V fails to be continuous at x, then $x = \bar{x}_j \in \mathcal{I}$ and $f_+(x,q) \geq 0$ for all $q \in \mathcal{Q}$. It therefore remains to show that $f_-^\phi(x) = (H_-)_p(x,p) \leq 0$.

Proposition C.1(vi) implies that $V^*(x) = \max\{V_-^{\mathrm{sc}}(x), V_{\mathcal{J}}^{\mathrm{sc}}(x)\}$. If $V^*(x) = V_{\mathcal{J}}^{\mathrm{sc}}(x)$, then according to (18), we have $\rho V_*(x) \geq H_j^{\mathcal{J}}(x) = \rho V_{\mathcal{J}}^{\mathrm{sc}}(x) = \rho V^*(x)$, and V is actually continuous at x, which is ruled out by hypothesis. So assume that $V^*(x) = V_-^{\mathrm{sc}}(x)$, then for all $p \in D^+ V^*(x) = (-\infty, p_-]$ we have that $H_-(x, p_-) = \rho V^*(x) \leq H_-(x,p)$, and consequently $f^\phi(x) = (H_-)_p(x, p_-) \leq 0$, which had to be proved.

Value at interface steady states. Condition (v) is a direct consequence of (18).

Strong push–push steady state. To show Condition (vi), let x be a strong push–push steady state. By hypothesis, we have $f_-^\phi > 0 > f_+^\phi$. Let $\lambda \in (0,1)$ be such that $\lambda f_-^\phi + (1-\lambda) f_+^\phi = 0$. Then $\rho V(x) = \lambda u_- + (1-\lambda) u_+$. We also have $\rho V(x) = H_-(p_-) = H_+(p_+)$. Combining these equalities, we see that

$$
\begin{aligned}
0 &= \lambda H_-(p_-) + (1-\lambda) H_+(p_+) - ((1-\lambda)u_+ + \lambda u_-) \\
&= \lambda p_- f_-^\phi + (1-\lambda) p_+ f_+^\phi = \lambda (p_- - p_+) f_-^\phi.
\end{aligned}
$$

As $\lambda \neq 0$ and $f_-^\phi \neq 0$, we infer that $p_- = p_+ = p^*$ and V_i^ϕ is differentiable at x, proving Condition (vi). This completes the proof of Theorem 2. □

F Proofs

F.1 Proof of Proposition A.1

Proof. Take $T > 0$ and $x \in \mathcal{X}$. According to the Dynamic Programming Principle (Bardi and Capuzzo-Dolcetta, 2008, Proposition III.2.5: although the assumptions are not fulfilled in our context, the proof carries over), we have

$$V(x) = \sup \left(\int_0^\theta \mathbf{u}(y(s), a(s)) \exp(-\rho s)\, \mathrm{d}s + V(y(\theta)) \exp(-\rho\theta) \right),$$

where $\theta = \min\{T, \Theta\}$ with Θ the exit time of y from $\mathcal{X}$, and where the supremum is taken over trajectory–action pairs (y, a) with $y(0) = x$.

Let (θ_k, y_k, a_k) be a sequence of time–trajectory–action triples with $y_k(0) = x$ such that

$$\int_0^{\theta_k} \mathbf{u}(y_k(s), a_k(s)) \exp(-\rho s)\, \mathrm{d}s + V(y_k(\theta_k)) \exp(-\rho\theta_k) \to V(x)$$

as $k \to \infty$, where $\theta_k = \min\{T, \Theta_k\}$ and Θ_k the exit time of y_k from $\mathcal{X}$. Introduce

$$w_k(t) = \int_0^{\min\{t,\theta_k\}} \mathbf{u}(y_k(s), a_k(s)) \exp(-\rho s)\, \mathrm{d}s + V(y_k(\theta_k)) \exp(-\rho\theta_k).$$

Then $w_k(T) \to V(x)$ as $k \to \infty$ and $\dot{w}_k(t)$ is measurable for $t \in [0, T]$. Extend (y_k, a_k) to $[0, T]$ by setting $(y_k(t), a_k(t)) = (y_k(\theta_k), a_k(\theta_k))$ if $\theta_k < t \leq T$. As θ_k is bounded, after restricting to a subsequence we may assume that $\theta_k \to \bar{\theta}$ as $k \to \infty$.

Introduce set-valued maps $\Phi_j : [0, T] \times \mathcal{X} \rightsquigarrow \mathbb{R}^2$ by setting

$$\Phi_j(t, z) = \{(\eta_0, \eta) \;:\; -\|u\|_\infty \leq \eta_0 \leq u_j(z, q) \exp(-\rho t),\ \eta = f_j(z, q),\ q \in \mathcal{Q}\}$$

if $(t, z) \in [0, T] \times \mathcal{X}_j$ and $\Phi_j(t, z) = \emptyset$ everywhere else. The sets $\Phi_j(t, z)$ are compact and convex by Assumption 2. Define $\Phi : [0, T] \times \mathcal{X} \rightsquigarrow \mathbb{R}^2$ by setting

$$\Phi(t, z) = \overline{\mathrm{co}} \left(\bigcup_{j=1}^{J} \Phi_j(t, z) \right).$$

Then $\Phi(t, z)$ is also compact and convex. Moreover, for all (t, z) it satisfies Property (Q) of Cesari (1983), that is,

$$\Phi(t, z) = \bigcap_{\delta > 0} \overline{\mathrm{co}} \left(\bigcup_{\|(\tilde{t}, \tilde{z}) - (t, z)\| < \delta} \Phi(\tilde{t}, \tilde{z}) \right).$$

We have that $(\dot{w}_k(t), \dot{y}_k(t)) \in \Phi(t, y_k(t))$ for all k and almost all $t \in [0, \theta_k]$, hence $|\dot{w}_k(t)| \leq \|u\|_\infty$ and $|\dot{y}_k(t)| \leq \|f\|_\infty$ almost everywhere on $[0, T]$. It follows that the (w_k, y_k) are equicontinuous, and a subsequence converges uniformly to a limit (w, y) on $[0, T]$. After relabelling, we may assume that the sequence itself converge uniformly to (w, y). By Cesari (1983, Theorem 8.6.i), it follows that $(\dot{w}(t), \dot{y}(t)) \in \Phi(t, y(t))$ for almost all $t \in [0, \bar{\theta}]$. Moreover, as the w_k converge uniformly, we have $w(T) = V(x)$. By the Filippov selection theorem (Vinter, 2000, Theorem 2.3.13), there is a measurable action schedule a such that $\dot{y}(t) = \mathbf{f}(y(t), a(t))$ almost everywhere on $[0, \bar{\theta}]$ and such that

$$V(x) = \int_0^{\bar{\theta}} \mathbf{u}(y(t), a(t)) \exp(-\rho t)\, \mathrm{d}t + V(y(\bar{\theta})) \exp(-\rho\bar{\theta}).$$

Set $y^*(t) = y(t)$ for $t \in [0, \bar{\theta}]$.

If $\bar{\theta} = T$, we repeat the argument with $x = y^*(T)$ and setting $(y^*(t), a^*(t)) = (y(t - T), a(t - T))$ for $t \in (T, 2T]$. Continuing inductively, we construct a trajectory–action pair (y^*, a^*) defined on an interval $[0, \Theta]$ such that $V(x) = U(y^*, a^*)$. □

F.2 Proof of Proposition A.3

Proof. Take $x, \tilde{x} \in \mathcal{X}$ such that $\tilde{x} < x$, and let a be such that $V(x) = U(y, a)$. Let $\Theta = \inf\{t : y(t) \notin \mathcal{X}\}$. Construct a real-valued function $\tilde{a}_0(t)$ such that $u(y(t), \tilde{a}_0(t)) = \mathbf{u}(y(t), a(t))$ for all $0 \le t \le \Theta$: that is, if $y(t) \in \mathcal{X} \backslash \mathcal{I}$, set $\tilde{a}_0(t) = \sum_{j=1}^{J} a_j(t) \mathbf{1}_{\mathcal{X}_j}(y(t))$; if $y(t) \in \mathcal{I}_j$, take $\tilde{a}_0(t)$ such that

$$u(y(t), \tilde{a}_0(t)) = \mu_j(a(t)) u(y(t), a_j(t)) + (1 - \mu_j(a(t))) u(y(t), a_{j+1}(t));$$

finally, if $t > \Theta$, set $\tilde{a}_0(t)$ to an arbitrary constant value in $\mathcal{Q}$.

Using $\tilde{a}_0$, we define an action schedule $\tilde{a}$ by setting $\tilde{a}_j(t) = \tilde{a}_0(t)$ for all t and all $j \in \{1, \ldots, J\}$. By Proposition 4.2 there is a state trajectory $\tilde{y}(t)$ such that $\tilde{y}(0) = \tilde{x}$ and $(\tilde{y}, \tilde{a})$ is a state–action pair.

Let $\tau = \min\{t : \tilde{y}(t) = y(t)\}$ and $\tilde{\Theta} = \inf\{t : \tilde{y}(t) \notin \mathcal{X}\}$. If $\tau \le \min\{\Theta, \tilde{\Theta}\}$, we have for all $0 \le t < \tau$ that $\tilde{y}(t) < y(t)$ and therefore $\mathbf{u}(\tilde{y}(t), \tilde{a}(t)) = u(\tilde{y}(t), \tilde{a}_0(t)) > u(y(t), \tilde{a}_0(t)) = \mathbf{u}(y(t), a(t))$, while for $t \ge \tau$, the trajectory-control pairs and their felicity flows are equal.

Take now $\tau > \min\{\Theta, \tilde{\Theta}\}$. For $\Theta < t < \tilde{\Theta}$, Assumption 3 implies that $\mathbf{u}(\tilde{y}(t), \tilde{a}(t)) > \rho\beta(y(\Theta))$, while for $\tilde{\Theta} < t < \Theta$, it implies $\rho\beta(\tilde{y}(\tilde{\Theta})) > \mathbf{u}(y(t), a(t))$. Finally, if $t \ge$

$\max\{\Theta, \tilde{\Theta}\}$, we have $\beta(\tilde{y}(\tilde{\Theta})) \geq \beta(y(\Theta))$. This proves the result. $\square$

F.3 Proof of Proposition A.4

Proof. Let $\sigma = \exp(-\rho t)\,\mathrm{d}t$ be the Borel measure on $[0, \infty)$ defined by $\sigma([t_1, t_2]) = (\exp(-\rho t_1) - \exp(-\rho t_2))/\rho$. The set $\{t \in \mathcal{T} : y(t) \neq x\}$ can be written as the union of at most countably many intervals $I_k = (t_{1,k}, t_{2,k})$ such that $\sigma(I_k) > 0$ and $y(t_{1,k}) = y(t_{2,k}) = x$, where K is the number of such intervals, and one interval $\hat{I} = (\hat{t}, \infty)$ such that $y(\hat{t}) = x$, which may be empty. Let $I_0 = [0, \infty) \backslash \left(\bigcup_{k=1}^{K} I_k \cup \hat{I}\right)$: this set is measurable, possibly of measure 0.

For $0 \leq k \leq K$ such that $\sigma(I_k) > 0$, introduce

$$v_k \equiv \frac{1}{\sigma(I_k)} \int_{I_k} \mathbf{u}(y(t), a(t)) \exp(-\rho t)\,\mathrm{d}t;$$

if $\sigma(\hat{I}) > 0$, set $\hat{v} \equiv (1/\sigma(\hat{I})) \int_{\hat{I}} w(t) \exp(-\rho t)\,\mathrm{d}t$ with $w(t) = \mathbf{u}(y(t), a(t))$ if $\hat{t} < t \leq \Theta$ and $w(t) = \beta(y(\Theta))$ if $t > \Theta$; finally $v_0 = 0$ if $\sigma(I_0) = 0$. Then

$$\bar{U} \equiv \int_0^{\Theta} \mathbf{u}(y(t), a(t)) \exp(-\rho t)\,\mathrm{d}t + \exp(-\rho\Theta)\beta(y(\Theta)) = \hat{v}\sigma(\hat{I}) + \sum_{k=0}^{K} v_k \sigma(I_k).$$

As $\sigma([0, \infty)) = 1/\rho$, either there exists $k \in \{0, \ldots\}$ such that $v_k \geq \rho\bar{U}$ and $\sigma(I_k) > 0$, or $\hat{v} \geq \rho\bar{U}$ and $\sigma(\hat{I}) > 0$.

Assume first that the first alternative holds for $k > 0$, and $y(t) > x$ for $t \in I_k$. Set $\Delta = t_{2,k} - t_{1,k} > 0$ and construct a trajectory-action pair by setting for $\ell = 0, 1, 2, \ldots$

$$(\tilde{y}(t), \tilde{a}(t)) = (y(t_{1,k} + t - \ell\Delta), a(t_{1,k} + t - \ell\Delta)), \qquad \text{if} \quad \ell\Delta \leq t < (\ell + 1)\Delta$$

We have

$$
\begin{aligned}
U(\tilde{y},\tilde{a}) &= \int_0^\infty \mathbf{u}(\tilde{y}(t),\tilde{a}(t))\exp(-\rho t)\,\mathrm{d}t \\
&= \sum_{\ell=0}^{\infty}\int_{\ell\Delta}^{(\ell+1)\Delta} \mathbf{u}(y(t_{1,k}+t-\ell\Delta),a(t_{1,k}+t-\ell\Delta))\exp(-\rho t)\,\mathrm{d}t \\
&= \exp(\rho t_{1,k})\int_{t_{1,k}}^{t_{2,k}} \mathbf{u}(y(s),a(s))\exp(-\rho s)\,\mathrm{d}s\sum_{\ell=0}^{\infty}\exp(-\rho\ell\Delta) \\
&= \frac{1-\exp(-\rho\Delta)}{\rho}v_k\frac{1}{1-\exp(-\rho\Delta)} = v_k/\rho \geq \bar{U}.
\end{aligned}
$$

Moreover $\tilde{y}(t) > x$ for almost all $t \geq 0$. Hence we have constructed the required trajectory. The argument for the situation that $y(t) < x$ for $t \in I_k$ is entirely analogous.

If the first alternative holds for $k = 0$, then the set C_0 of constant actions q such that $\mathbf{f}(x, q) = 0$ is non-empty. As C_0 is compact, there is a maximiser $\bar{q}$ of $\mathbf{u}(x, q)$ restricted to C_0. Let $(\tilde{y}, \tilde{a})$ be the trajectory–action pair $\tilde{y}(t) = x$, $\tilde{a}(t) = \bar{q}$ for all t. Then $\mathbf{u}(y, a) \leq \mathbf{u}(\tilde{y}, \tilde{a})$ for all $t \in I_0$, and

$$
\begin{aligned}
\rho\bar{U} \leq v_0 &= \frac{1}{\sigma(I_0)}\int_{I_0}\mathbf{u}(\tilde{y}(t),\tilde{a}(t))\exp(-\rho t)\,\mathrm{d}t = \frac{1}{\sigma(I_0)}\int_{I_0}\mathbf{u}(x,\bar{q})\exp(-\rho t)\,\mathrm{d}t \\
&= \mathbf{u}(x,\bar{q}) = \rho U(\tilde{y},\tilde{a}),
\end{aligned}
$$

completing the construction of the trajectory also in this situation.

Finally, if $\hat{v} \geq \rho\bar{U}$ and $\sigma(\hat{I}) > 0$, then $(\tilde{y}(t), \tilde{a}(t)) = (y(\hat{t} + t), a(\hat{t} + t))$ achieves a higher payoff than $\bar{U}$. □

F.4 Proof of Proposition A.5

Proof. Let $\Delta = \min_{j\neq k} |\bar{x}_j - \bar{x}_k|$, and let $M = \|\mathbf{f}\|_\infty > 0$. Introduce for a trajectory–action pair $\pi = (y, a)$ the exit time $\Theta(\pi) = \inf\{t \geq 0 : y(t) \notin \mathcal{X}\}$, the time interval $\mathcal{T}(\pi) = [0, \Theta(\pi)]$, and the set $S_j(\pi)$ of singular pull-pull events as

$$
S_j(\pi) \equiv \{t \in \mathcal{T}(\pi) \,:\, y(t) = \bar{x}_j,\ f_j(y(t), a_j(t)) < 0\,,\ f_{j+1}(y(t), a_{j+1}(t)) > 0\}\,.
$$

If π is not regular, the union $\bigcup_j S_j(\pi)$ has positive Lebesgue measure.

Let π be a given trajectory–control pair. For $\ell = 1, 2, \ldots$, we shall inductively construct a sequence $\pi^{(\ell)} = (y^{(\ell)}, a^{(\ell)})$ of trajectory–action pairs such that $\pi^{(0)} = \pi$, $S_j(\pi^{(\ell)}) \cap [0, \ell\Delta/M)$ has measure zero for every j, and $U(\pi^{(\ell+1)}) \geq U(\pi^{(\ell)})$ for all $\ell \geq 0$.

Assume that $\pi^{(\ell)}$ has already been constructed. Let

$$\tau = \inf\left\{t_1 \in \mathfrak{T}(\pi^{(\ell)}) \,:\, S_j(\pi^{(\ell)}) \cap [0, t_1] \text{ has positive measure for some } j\right\}.$$

If $\tau \geq (\ell+1)M/\Delta$, then we set $\pi^{(\ell+1)} = \pi^{(\ell)}$ and the induction step is completed.

If $\tau < (\ell+1)M/\Delta$, then $y^{(\ell)}(\tau) = \bar{x}_j$ for some j, and we set $\pi_\tau(t) = \pi^{(\ell)}(t - \tau)$.

By Proposition A.4, there is a trajectory–action pair $\tilde{\pi} = (\tilde{y}, \tilde{a})$ such that either $\tilde{y}(t) < \bar{x}_j$ for all $t \geq 0$, or $\tilde{y}(t) > \bar{x}_j$ for all $t \geq 0$, or $\tilde{y}(t) = \bar{x}_j$ for all $t \geq 0$, as well as $U(\tilde{\pi}) \geq U(\pi_\tau)$. In the first two cases, $\tilde{y}(t) \notin \mathfrak{I}\backslash\mathfrak{I}_j$ for all $0 \leq t < \Delta/M$, as for those values of t we have $|y(t) - y(0)| \leq Mt < \Delta$. In these situations we set $\pi^{(\ell+1)}(t) = \pi^{(\ell)}(t)$ for $0 \leq t \leq \tau$ and $\pi^{(\ell+1)}(t) = \tilde{\pi}(t - \tau)$ for $t \geq \tau$. Then

$$\begin{aligned}
U(\pi^{(\ell+1)}) &= \int_0^\tau \mathbf{u}(y^{(\ell)}(t), a^{(\ell)}(t)) \exp(-\rho t)\,\mathrm{d}t + \exp(-\rho\tau)U(\tilde{\pi}) \\
&\geq \int_0^\tau \mathbf{u}(y^{(\ell)}(t), a^{(\ell)}(t)) \exp(-\rho t)\,\mathrm{d}t + \exp(-\rho\tau)U(\pi^{(\ell)}) = U(\pi^{(\ell)}).
\end{aligned}$$

In the third case, according to Proposition A.4, we may assume that $\tilde{\pi}$ is generated by a constant action schedule $\tilde{a}(t) = q$ for all $t \geq 0$. If $\tilde{\pi}$ is a regular trajectory, then we define $\pi^{(\ell+1)}$ as in the first two cases. If $\tilde{\pi}$ is singular, then in particular $f_j(\bar{x}_j, q_j) < 0$. Consider the trajectory–action pair $(z, \tilde{a})$ that satisfies $z(0) = \bar{x}_j$ and $\dot{z}(t) = f_j(z(t), q_j)$ for $0 \leq t < M/\Delta$. As before, we have that $z(t) \notin \mathfrak{I}$ for $0 < t < M/\Delta$ and, as $f_j(\bar{x}_j, q_j) < 0$, we also have that $z(t) < \bar{x}_j = y(t)$ for all $t > 0$. By Proposition A.2, it follows that $U(z, \tilde{a}) > U(y, \tilde{a})$. Setting $\pi^{(\ell+1)}(t) = \pi^{(\ell)}(t)$ for $0 \leq t \leq \tau$ and $\pi^{(\ell+1)}(t) = (z(t-\tau), \tilde{a}(t-\tau))$ for $t \geq \tau$, and noting that also in this case $U(\pi^{(\ell+1)}) \geq U(\pi^{(\ell)})$ finishes the inductive step.

The induction either breaks off at the ℓ'th step and produces a regular trajectory, as indicated, or it continuous indefinitely. In the latter case, we set $\bar{\pi}(t) = \lim_{\ell\to\infty} \pi^{(\ell)}(t)$. Then $S_j(\bar{\pi})$ has measure zero for all j and $\bar{y}$ is regular also in this case. $\square$

F.5 Proof of Proposition C.1

Proof of Proposition C.1. Throughout the proof, we write $\bar{x}$ for $\bar{x}_j$, f_- for f_j and f_+ for f_{j+1} etc. In particular $\mathcal{X}_+$ denotes $\mathcal{X}_{j+1}$ and not $\mathcal{X}_{j,+}$. We fix $\bar{q}$ such that $u(\bar{x}, \bar{q}) \geq u(\bar{x}, q)$ for all $q \in \mathcal{Q}$.

Statement (i) is a direct corollary of Proposition A.4.

To prove (ii), assume that the dynamics are right controllable at $\bar{x}$: the argument for left controllability is analogous.

By controllability and continuity of f_+, there are $\delta, m > 0$ such that $[\bar{x}, \bar{x} + \delta] \in \mathcal{X}_+$ and $[-m, m] \subset f_+(\bar{x}, \mathcal{Q})$ for all $\bar{x} \leq x \leq \bar{x} + \delta$. Take $x_1, x_2 \in [\bar{x}, \bar{x} + \delta]$ as well as $\sigma \in \{-m, m\}$ such that $y(t) = x_1 + \sigma t$ satisfies $y(0) = x_1$ and $y(\tau) = x_2$ if $\tau = |x_2 - x_1|/m$.

As $|\dot{y}(t)| = m$ and $y(t) \in [\bar{x}, \bar{x} + \delta]$ for $0 \leq t \leq \tau$ there is $a(t)$ such that $\dot{y}(t) = f(y(t), a(t))$ for all $0 \leq t \leq \tau$. Then

$$V(x_1) \geq \int_0^\tau u_+(y(t), a(t)) \exp(-\rho t) \, \mathrm{d}t + V(x_2) \exp(-\rho \tau).$$

As $|V(x)| \leq \|\mathbf{u}\|_\infty / \rho$ for all x, we obtain

$$V(x_1) - V(x_2) \geq -|\exp(-\rho\tau) - 1| \|\mathbf{u}\|_\infty / \rho - \|\mathbf{u}\|_\infty \tau \geq -2\|\mathbf{u}\|_\infty \tau.$$

Interchanging the roles of x_1 and x_2, and using the definition of τ, then gives

$$|V(x_1) - V(x_2)| \leq \frac{2\|\mathbf{u}\|_\infty}{m} |x_1 - x_2|.$$

For (iii), let $\bar{x}$ be a right semi-attractor: then $f_+(\bar{x}, q) \leq 0$ for all $q \in \mathcal{Q}$.

Choosing (y, c) such that $y(0) = \bar{x}$ and $a_-(t) = a_+(t) = \bar{q}$ for all $t \in \mathcal{T}$ implies first that $y(t) \leq \bar{x}$ for all $t \in \mathcal{T}$, as $\bar{x}$ is a right semi-attractor. Since $u_x(x, \bar{q}) < 0$ for all x, we then have $u(y(t), \bar{q}) \geq u(\bar{x}, \bar{q})$ for all $t \geq 0$ and hence $V(\bar{x}) \geq u(\bar{x}, \bar{q})/\rho$.

Take $x \geq \bar{x}$, and let now the pair (y, c) be such that $y(0) = x$ and $V(x) = U(y, c)$.

Introduce $\theta = \inf\left\{t \in \mathcal{T} \,:\, y(t) \notin \mathring{\mathcal{X}}_{j,+}\right\}$. We have

$$\begin{aligned}V(x) &= \int_0^{\min\{\theta,\Theta\}} u(y(t), a_+(t)) \exp(-\rho t)\,\mathrm{d}t \\ &\qquad + V(\bar{x}) \exp(-\rho\theta) \mathbf{1}_{\{t\,:\,t\leq\Theta\}}(\theta) + V(y(\Theta)) \exp(-\rho\Theta) \mathbf{1}_{\{t\,:\,t>\Theta\}}(\theta) \\ &\leq \int_0^{\min\{\theta,\Theta\}} u(\bar{x}, a_+(t)) \exp(-\rho t)\,\mathrm{d}t \\ &\qquad + V(\bar{x}) \exp(-\rho\theta) \mathbf{1}_{\{t\,:\,t\leq\Theta\}}(\theta) + V(y(\Theta)) \exp(-\rho\Theta) \mathbf{1}_{\{t\,:\,t>\Theta\}}(\theta) \\ &\leq \int_0^{\min\{\theta,\Theta\}} u(\bar{x}, \bar{q}) \exp(-\rho t)\,\mathrm{d}t + V(\bar{x}) \exp(-\rho \min\{\theta,\Theta\}) \leq V(\bar{x}).\end{aligned}$$

This shows right upper semi-continuity of V at $\bar{x}$.

Proposition IV.3.4 of Bardi and Capuzzo-Dolcetta (2008) implies that the value function is also lower semi-continuous at $\bar{x}$. This then establishes right continuity.

Finally, if there is a trajectory starts at $\bar{x}$ and remains in $\mathcal{X}_{j,+}$ for all $t \geq 0$, it must be equal to $y(t) = \bar{x}$. Hence $V_+^{\mathrm{sc}}(\bar{x}) \leq V_{\mathcal{J}}^{\mathrm{sc}}(\bar{x})$, which shows the second part of the statement.

For (iv), let $\bar{x}$ be a left semi-attractor.

If $f_-(\bar{x}, q_\ell) > 0$, then there are $m > 0$ and $\delta > 0$ such that $f_-(z, q) > m$ for all $z \in [\bar{x}-\delta, \bar{x}]$ and all $q \in \mathcal{Q}$. Fix $x \in [\bar{x} - \delta, \bar{x}]$, and let (y, a) be a trajectory–action pair such that $y(0) = x$ and $V(x) = U(y, a)$. Then there is $0 < \tau < |x - \bar{x}|/m$ such that $y(t) < \bar{x}$ for all $0 < t < \tau$ and $y(\tau) = \bar{x}$. This implies

$$\begin{aligned}V(x) - V(\bar{x}) &= \left(\int_0^\tau \mathbf{u}(y(t), a(t)) \exp(-\rho t)\,\mathrm{d}t + (\exp(-\rho\tau) - 1) V(\bar{x})\right) \\ &\leq 2\|u\|_\infty \tau = \frac{2\|u\|_\infty}{m} |x - \bar{x}|,\end{aligned}$$

which shows that V is left upper semi-continuous at $\bar{x}$. Proposition IV.3.4 of Bardi and Capuzzo-Dolcetta (2008) again ensures left continuity.

If $f_-(\bar{x}, q_\ell) = 0$, take $\varepsilon > 0$ and $\delta > 0$. Let T be the unique solution of $\exp(-\rho T)\|u\|_\infty/\rho = \varepsilon/2$; this solution is positive if $\varepsilon > 0$ is sufficiently small. Let $L_f > 0$ be such that $|f_-(z, q_\ell)| \leq L_f |z - \bar{x}|$ for all $-\delta < z - \bar{x} < 0$, and let $\delta_1 = \exp(-L_f T)\delta$.

Take $x \in (\bar{x} - \delta_1, \bar{x})$, and let y be any state trajectory with $y(0) = x$. Set $\tau = \inf\{t \geq 0 \,:\, y(t) \in \mathcal{X}_+\}$, and let $\theta = \min\{\tau, T\}$. Using the Gronwall inequality, we have that

$-\delta < -\exp(L_f t)|x-\bar{x}| \le y(t) - \bar{x} \le 0$ for all $0 \le t \le \theta$.

Let now (y, a) be a trajectory–action pair such that $y(0) = x$ and $V(x) = U(y, a)$. To obtain an estimate for the payoff on the time interval $[0, \theta]$, we split it in a part $\mathcal{T}_1$ where the state moves quickly to the right, which restricts the amount of time it can spend in this set, and a part $\mathcal{T}_2 \cup \mathcal{T}_3$ where it moves slowly to the right, or not at all, restricting the value of $a_-(t)$ from above, and hence the payoff.

Take $\eta > 0$ and form the partition $\mathcal{T}_1 \cup \mathcal{T}_2 \cup \mathcal{T}_3 \cup \mathcal{T}_4$ of the interval $[0, \theta]$, where $\mathcal{T}_1 = \{t : \dot{y}(t) > \eta\}$, $\mathcal{T}_2 = \{t : 0 \le \dot{y}(t) \le \eta\}$, $\mathcal{T}_3 = \{t : \dot{y}(t) < 0\}$, and $\mathcal{T}_4 = \{t : y \text{ is not differentiable at } t\}$. Note that $\mathcal{T}_4$ is a set of measure zero.

Clearly

$$y(\theta) - x = \int_0^\theta \dot{y}(t)\,\mathrm{d}t = \int_{\mathcal{T}_1} + \int_{\mathcal{T}_2} + \int_{\mathcal{T}_3} \dot{y}(t)\,\mathrm{d}t.$$

Since $-\exp(L_f\theta)|x-\bar{x}| \le y(\theta) - \bar{x} \le 0$, the measure $|\mathcal{T}_1|$ of the first partitioning set satisfies

$$\begin{aligned}
\eta|\mathcal{T}_1| &\le \int_{\mathcal{T}_1} \dot{y}(t)\,\mathrm{d}t \le \int_{\mathcal{T}_1} \dot{y}(t)\,\mathrm{d}t + \int_{\mathcal{T}_2} \dot{y}(t)\,\mathrm{d}t = y(\theta) - x - \int_{\mathcal{T}_3} \dot{y}(t)\,\mathrm{d}t \\
&\le |y(\theta) - x| - \int_{\mathcal{T}_3} \dot{y}(t)\,\mathrm{d}t \le |y(\theta) - \bar{x}| + |\bar{x} - x| - \int_{\mathcal{T}_3} f_-(y(t), q_\ell)\,\mathrm{d}t \\
&\le (1 + \exp(L_f\theta))|x - \bar{x}| + \int_{\mathcal{T}_3} L_f|y(t) - \bar{x}|\,\mathrm{d}t \\
&\le \left(1 + \exp(L_f\theta) + L_f\theta\exp(L_f\theta)\right)|x - \bar{x}| =: C_1|x - \bar{x}|.
\end{aligned}$$

Consequently, the integral of the discounted flow payoff evaluated over $\mathcal{T}_1$ is bounded by

$$\int_{\mathcal{T}_1} u(y(t), a(t))\exp(-\rho t)\,\mathrm{d}t \le \|u\|_\infty|\mathcal{T}_1| \le \frac{C_1\|u\|_\infty}{\eta}|x - \bar{x}|.$$

To estimate the payoff evaluated over $\mathcal{T}_2 \cup \mathcal{T}_3$, we need an upper bound on $a_-(t)$. Let $L_u, \ell_f > 0$ be such that $|u(z, q) - u(\bar{x}, q_\ell)| \le L_u(|z - \bar{x}| + |q - q_\ell|)$ for all (z, q), and $0 < \ell_f < \frac{\partial f_-}{\partial q}(\bar{x}, q)$ for all $q \in \mathcal{Q}$: such constants exist as a consequence of Assumptions 1 and 2 and the compactness of $\mathcal{Q}$. We have for $t \in \mathcal{T}_2 \cup \mathcal{T}_3$ that

$$\begin{aligned}\eta \geq \dot{y}(t) &= f_-(y(t), a_-(t)) \geq f_-(y(t), q_\ell) + \ell_f(a_-(t) - q_\ell)\\ &\geq -L_f|y(t) - \bar{x}| + \ell_f(a_-(t) - q_\ell);\end{aligned}$$

in the last inequality we used that $f(\bar{x}, q_\ell) = 0$. Hence

$$a_-(t) - q_\ell \leq (\eta/\ell_f) + (L_f/\ell_f)|y(t) - \bar{x}|$$

and

$$\begin{aligned}|u(y(t), a_-(t)) - u(\bar{x}, q_\ell)| &\leq L_u \left(|y(t) - \bar{x}| + |a_-(t) - q_\ell|\right)\\ &\leq L_u\left(1 + \frac{L_f}{\ell_f}\right)|y(t) - \bar{x}| + \frac{L_u}{\ell_f}\eta.\end{aligned}$$

This implies

$$\begin{aligned}&\int_{\mathcal{T}_2 \cup \mathcal{T}_3} u(y(t), a(t)) \exp(-\rho t)\,\mathrm{d}t\\ &\leq \int_0^\theta \Big(u(\bar{x}, q_\ell) + L_u(1 + L_f/\ell_f)|y(t) - \bar{x}| + L_u\eta/\ell_f\Big) \exp(-\rho t)\,\mathrm{d}t\\ &\leq (1 - \exp(-\rho\theta))u(\bar{x}, q_\ell)/\rho + C_2|x - \bar{x}| + C_3\eta,\end{aligned}$$

where $C_2 = L_u(1 + L_f/\ell_f)\exp(L_f T)$ and $C_3 = T L_u/\ell_f$.

Combining these estimates yields

$$\begin{aligned}V(x) &\leq \int_{\mathcal{T}_1} + \int_{\mathcal{T}_2} + \int_{\mathcal{T}_3} u(y(t), a(t))\,\mathrm{d}t + \exp(-\rho\theta)V(y(\theta))\\ &\leq (1 - \exp(-\rho\theta))u(\bar{x}, q_\ell)/\rho + \exp(-\rho\theta)V(y(\theta))\\ &\quad + \frac{C_1\|u\|_\infty}{\eta}|x - \bar{x}| + C_2|x - \bar{x}| + C_3\eta.\end{aligned}$$

Choose $\eta = \varepsilon/(6C_3)$ and $|x - \bar{x}| < \min\{\delta_1, \varepsilon^2/(36C_1C_3\|u\|_\infty), \varepsilon/(6C_2)\}$, and recall that $V_-^{\mathrm{sc}}(\bar{x}) = u(\bar{x}, q_\ell)/\rho \leq V(\bar{x})$, to obtain

$$V(x) \leq (1 - \exp(-\rho\theta))V(\bar{x}) + \exp(-\rho\theta)V(y(\theta)) + \varepsilon/2.$$

If $\theta = T$, then $\exp(-\rho\theta)V(y(\theta)) \leq \varepsilon/2$ and $V(x) \leq V(\bar{x}) + \varepsilon$, showing that V is left upper semi-continuous at $\bar{x}$. If $\theta = \tau$, then $V(y(\theta)) = V(\bar{x})$ and $V(x) \leq V(\bar{x}) + \varepsilon/2$, again showing that V is left upper semi-continuous at $\bar{x}$. As lower semi-continuity is assured, it follows that V is left continuous at $\bar{x}$.

A trajectory starting at $\bar{x}$ and remaining in $\mathcal{X}_-$ for all $t \geq 0$ must satisfy $y(t) = \bar{x}$ for all t: therefore $V_-^{\mathrm{sc}}(\bar{x}) \leq V_{\mathrm{J}}^{\mathrm{sc}}(\bar{x})$.

Proceeding to (v), let $\bar{x}$ be a left semi-repeller, and let (y, a) be a trajectory–action pair such that $y(0) = \bar{x}$, $a_-(t) = a_+(t) = \bar{q}$ and $y(t) \in \mathcal{X}_-$ for all $t \geq 0$. Then $V_-^{\mathrm{sc}}(\bar{x}) \geq U(y, a) \geq u(\bar{x}, \bar{q})/\rho \geq V_{\mathrm{J}}^{\mathrm{sc}}(\bar{x})$. Let $(\tilde{y}, \tilde{a})$ be any trajectory–action pair such that $\tilde{y}(0) = \bar{x}$ and $\tilde{y}(t) \in \mathcal{X}_+$ for all $t \geq 0$. Then

$$u(\bar{x}, \bar{q}) - u(\tilde{y}(t), \tilde{a}_+(t)) \geq u(\bar{x}, \bar{q}) - u(\bar{x}, \tilde{a}_+(t)) + u(\bar{x}, \tilde{a}_+(t)) - u(\tilde{y}(t), \tilde{a}_+(t)) \geq 0,$$

as $\bar{q}$ maximises $u(\bar{x}, \cdot)$ and as $u(x, q)$ is decreasing in x. This implies $V_-^{\mathrm{sc}}(\bar{x}) \geq u(\bar{x}, \bar{q})/\rho \geq V_+^{\mathrm{sc}}(\bar{x})$. We conclude that $V(\bar{x}) = V_-^{\mathrm{sc}}(\bar{x})$, and hence that V is left continuous at $\bar{x}$.

Finally, we prove (vi) The sufficiency of the condition is clear. To show necessity, we combine (ii), (iv) and (v) to infer that V is always left-continuous. Statements (ii) and (iii) imply that it can only fail to be right-continuous if $\bar{x}$ is a right semi-repeller and $V_+^{\mathrm{sc}}(\bar{x}) < V(\bar{x})$. $\square$

F.6 Proof of Proposition C.3

First, we formulate superoptimality and suboptimality principles at interfaces.

Lemma F.1. *Let $\bar{x} = \bar{x}_j$ be an interface point, $v : \mathcal{X} \to \mathbb{R}$ a supersolution and $w : \mathcal{X} \to \mathbb{R}$ a subsolution of* (16)*, such that $v, w \in \mathscr{G}$. Choose $\xi_1 \in \mathring{\mathcal{X}}_j$, and let $\tau_j = \inf\{t \geq 0 : y(t) \notin (\xi_1, \bar{x})\}$ be the exit time from $(\xi_1, \bar{x})$. Then for all $t \geq 0$ and all $x \in (\xi_1, \bar{x})$, we have for $\theta_j = \min\{t, \tau_j\}$ that*

$$v(x) \geq \sup_a \left(\int_0^{\theta_j} \mathbf{u}(y(s), a(s)) \exp(-\rho s) \, \mathrm{d}s + \exp(-\rho\theta_j) v(y(\theta_j)) \right) \tag{25}$$

and

$$w(x) \leq \sup_a \left(\int_0^{\theta_j} \mathbf{u}(y(s), a(s)) \exp(-\rho s) \, \mathrm{d}s + \exp(-\rho\theta_j) w(y(\theta_j)) \right). \tag{26}$$

If v or w are, respectively, continuous at $\bar{x}$, $\xi_2 \in \mathring{\mathcal{X}}_{j+1}$ and $\tau_{j+1} = \inf\{t \geq 0 : y(t) \notin (\bar{x}, \xi_2)\}$, then for all $t \geq 0$ and $x \in (\bar{x}, \xi_2)$, the inequalities (25) *or* (26) *hold, respectively, with θ_j replaced by $\theta_{j+1} = \min\{t, \tau_{j+1}\}$.*

Proof. As the hypotheses imply that v and w restricted to $[\xi_1, \bar{x}]$ and $[\bar{x}, \xi_2]$ are continuous, equation (25) is implied by Bardi and Capuzzo-Dolcetta (2008, Remark III.2.34), and (26) by Bardi and Capuzzo-Dolcetta (2008, Remark IV.3.16). □

Proof of Proposition C.3. Write H_- for H_j and H_+ for H_{j+1}; then we have that $\mathbf{H}(\bar{x}, p) = \max\{H_-(\bar{x}, p), H_+(\bar{x}, p)\}$. Introduce $\varphi(p) = \rho w(\bar{x}) - \mathbf{H}(\bar{x}, p)$.

Examine first the situation that $\varphi(p) \leq 0$ for all p. As $\varphi(p)$ is an affine function of p if $p < p_\ell(\bar{x})$ or $p > p_u(\bar{x})$, a maximiser $\bar{p}$ of φ exists. The function φ is concave and maximal at $\bar{p}$, hence 0 is an element of the subgradient of $-\varphi(\bar{p})$, which is the closed convex hull of the derivatives $(H_-)_p(\bar{x}, \bar{p})$ and $(H_+)_p(\bar{x}, \bar{p})$ (Aubin, 1993, Corollary 4.4). Using the fact that $(H_\pm)_p(x, p) = f_\pm(x, q^*(x, p))$, and setting $\bar{q}^* = q^*(\bar{x}, \bar{p})$, this implies that there is $0 \leq \lambda \leq 1$ such that

$$\mathbf{H}(\bar{x}, \bar{p}) = \lambda H_-(\bar{x}, \bar{p}) + (1 - \lambda) H_+(\bar{x}, \bar{p}), \qquad \lambda f_-(\bar{x}, \bar{q}^*) + (1 - \lambda) f_+(\bar{x}, \bar{q}^*) = 0, \tag{27}$$

$\lambda = \mu(\bar{q}^*, \ldots, \bar{q}^*)$ and $(\bar{q}^*, \bar{q}^*) \in C_0$, where C_0 is the set of controls stabilising $\bar{x}$. Using (27), as well as the definition of $H^{\mathcal{I}}$, we obtain

$$\begin{aligned} 0 \geq \varphi(\bar{p}) &= \rho w(\bar{x}) - (\lambda H_-(\bar{x}, \bar{p}) + (1 - \lambda) H_+(\bar{x}, \bar{p})) \\ &= \rho w(\bar{x}) - \lambda u(\bar{x}, \bar{q}^*) - (1 - \lambda) u(\bar{x}, \bar{q}^*) - \bar{p}(\lambda f_-(\bar{x}, \bar{q}^*) + (1 - \lambda) f_+(\bar{x}, \bar{q}^*)) \\ &\geq \rho w(\bar{x}) - H^{\mathcal{I}}(\bar{x}). \end{aligned}$$

In this case the alternative A holds true.

Consider now the second situation, that there is $\bar{p}$ such that $\varphi(\bar{p}) > 0$. Let $\varepsilon > 0$ and set

$$\psi_\varepsilon(x) = w(\bar{x}) + \bar{p}(x - \bar{x}) + \frac{(x - \bar{x})^2}{2\varepsilon^2}.$$

Now $\bar{x}$ cannot maximise $w - \psi_\varepsilon$ for any $\varepsilon > 0$, for if it did, $\psi'_\varepsilon(\bar{x}) = \bar{p} \in D^+w(\bar{x})$, which would imply, as w is a subsolution, that $\varphi(\bar{p}) \leq 0$.

For every $\varepsilon > 0$ let x_ε denote a maximiser of $w - \psi_\varepsilon$. Necessarily $x_\varepsilon \neq \bar{x}$ and $0 = w(\bar{x}) - \psi_\varepsilon(\bar{x}) \leq w(x_\varepsilon) - \psi_\varepsilon(x_\varepsilon)$, which implies, first, with $\sigma = (x_\varepsilon - \bar{x})/|x_\varepsilon - \bar{x}|$, that

$$(x_\varepsilon - \bar{x})^2 + 2\sigma\varepsilon^2\bar{p}|x_\varepsilon - \bar{x}| \leq 2\varepsilon^2(w(x_\varepsilon) - w(\bar{x})) \leq 4\varepsilon^2\|w\|_\infty;$$

then $(|x_\varepsilon - \bar{x}| + \sigma\varepsilon^2\bar{p})^2 \leq \varepsilon^2\left(4\|w\|_\infty + \varepsilon^2\bar{p}^2\right)$; and finally $|x_\varepsilon - \bar{x}| \leq C\varepsilon$, where $C = (4\|w\|_\infty + \varepsilon^2\bar{p}^2)^{\frac{1}{2}} - \varepsilon\sigma\bar{p}$. So $x_\varepsilon \to \bar{x}$ as $\varepsilon \to 0$. In particular, if $\varepsilon > 0$ is sufficiently small, x_ε is in a neighbourhood of $\bar{x}$ containing only a single interface point, namely $\bar{x}$.

We can say more about x_ε if w is discontinuous at $\bar{x}$. As w is left continuous and non-increasing, there is $\zeta > 0$ such that $w(x) \leq w(\bar{x}) - \zeta$ if $x > \bar{x}$. Since w is non-increasing, for $\bar{x} < x \leq \bar{x} + C\varepsilon$, with C defined as above, we have that

$$w(x) - \psi_\varepsilon(x) \leq -\zeta - \bar{p}(x - \bar{x}) - \frac{(x - \bar{x})^2}{2\varepsilon^2} \leq -\zeta + C|\bar{p}|\varepsilon < 0$$

if $\varepsilon > 0$ is sufficiently small. Since $w(x_\varepsilon) - \psi_\varepsilon(x_\varepsilon) \geq 0$, it follows that $x_\varepsilon \leq \bar{x}$ if $\varepsilon > 0$ is sufficiently small.

We select a sequence $\varepsilon_k > 0$ such that $\varepsilon_k \to 0$, an index $\ell \in \{j, j+1\}$, and a sequence of maximisers x_k of $w - \psi_{\varepsilon_k}$ such that $x_k \in \mathfrak{X}_\ell$ for all k; by the previous remark, $x_k \in \mathfrak{X}_j$ for all k if w is discontinuous at $\bar{x}$. Actually, we can pick $\xi_1 \in \mathring{\mathfrak{X}}_j$ and $\xi_2 \in \mathring{\mathfrak{X}}_{j+1}$ such that either $x_k \in I_j \equiv (\xi_1, \bar{x})$ or $x_k \in I_{j+1} \equiv (\bar{x}, \xi_2)$ for all k, if necessary by discarding a finite number of the initial x_k.

Let $t > 0$ be sufficiently small such that any trajectory y starting at x_k satisfies $\xi_1 < y(s) < \xi_2$ for all $0 \leq s \leq t$. Let $\tau_\ell(y) = \inf\{s \geq 0 : y(s) \notin I_\ell\}$. By Lemma F.1, we have for $\theta_\ell(y) = \min\{t, \tau_\ell(y)\}$ that

$$w(x_k) \leq \sup_a \left(\int_0^{\theta_\ell(y)} u(y(s), a(s)) \exp(-\rho s)\, \mathrm{d}s + \exp(-\rho\theta_\ell) w(y(\theta_\ell)) \right). \tag{28}$$

For every k, let (y_k, a_k) be a trajectory–action pair starting at x_k that realises the supremum on the right hand side of (28), and let $\theta_\ell^{(k)} = \theta_\ell(y_k)$.

If the alternative B holds, we are done. So assume that it does not hold. Then $\theta_\ell^{(k)} \to 0$ as $k \to \infty$ and $\min\{t, \theta_\ell^{(k)}\} = \theta_\ell^{(k)}$ for k sufficiently large.

Note that $y_k(\theta_\ell^{(k)}) = \bar{x}$. From the fact that $w(x_\varepsilon) - \psi_\varepsilon(x_\varepsilon) \geq 0$, we derive $w(x_k) \geq \psi_{\varepsilon_k}(x_k) \geq w(\bar{x}) + \bar{p}(x_k - \bar{x})$. Combining this with (28) then yields

$$
\begin{aligned}
0 &\leq \int_0^{\theta_\ell^{(k)}} u(y_k(s), a_k(s)) \exp(-\rho s)\,\mathrm{d}s + (\exp(-\rho\theta_\ell^{(k)}) - 1)w(\bar{x}) - \bar{p}(x_k - \bar{x}) \\
&= \int_0^{\theta_\ell^{(k)}} u(y_k(s), a_k(s)) \exp(-\rho s)\,\mathrm{d}s + (\exp(-\rho\theta_\ell^{(k)}) - 1)w(\bar{x}) \\
&\qquad + \bar{p} \int_0^{\theta_\ell^{(k)}} f(y_k(s), a_k(s))\,\mathrm{d}s \\
&\leq \int_0^{\theta_\ell^{(k)}} \max_{q \in \mathcal{Q}} \left[u(y_k(s), q) \exp(-\rho s) + \bar{p} f(y_k(s), q)\right]\,\mathrm{d}s + (\exp(-\rho\theta_\ell^{(k)}) - 1)w(\bar{x}).
\end{aligned}
$$

Dividing by $\theta_\ell^{(k)}$ and taking the limit $k \to \infty$ then yields $0 \leq H_\ell(\bar{x}, \bar{p}) - \rho w(\bar{x})$, implying that $\varphi(\bar{p}) = \rho w(\bar{x}) - \max\{H_j(\bar{x}, \bar{p}), H_{j+1}(\bar{x}, \bar{p})\} \leq 0$, contradicting the choice of $\bar{p}$. $\square$

F.7 Proof of Proposition C.4

Proof. The continuous function $w - v$ takes on the compact set $\mathcal{X}$ a maximum M at a point $\bar{x}$. Assume that $M > 0$, as otherwise the lemma is proved.

If $\bar{x}$ is neither an interface point nor a boundary point of $\mathcal{X}$, the proof uses the classical "doubling of variables" technique, (see Bardi and Capuzzo-Dolcetta, 2008, Theorem II.3.1) to derive a contradiction.

If $\bar{x} \in \partial\mathcal{X}$, say $\bar{x} = \bar{x}_0$, the case $\bar{x} = \bar{x}_J$ being similar, then (17) and (19) imply either $w(\bar{x}) \leq \beta(\bar{x}) \leq v(\bar{x})$, contradicting $M > 0$, or that one of the following holds: $\rho w(\bar{x}) - H_1(\bar{x}, p) \leq 0$ for all $p \in D^+ w(\bar{x})$, or $\rho v(\bar{x}) - H_1(\bar{x}, p) \geq 0$ for all $p \in D^- w(\bar{x})$. The argumentation proceeds then as in the proof of Bardi and Capuzzo-Dolcetta (2008, Theorem V.4.16).

Hence we only have to consider the situation that $\bar{x}$ is an interface point. According to Proposition C.3, one of two alternatives can obtain. If Alternative A is true, then we have $\rho w(\bar{x}) \leq H^{\mathcal{I}}(\bar{x}) \leq \rho v(\bar{x})$, as the second inequality is implied by (18). This implies that $w(\bar{x}) - v(\bar{x}) = M \leq 0$, a contradiction.

If Alternative B holds, there is $\eta > 0$, $\ell \in \{j, j+1\}$, and a sequence $x_k \to \bar{x}$, such that $x_k \in \mathcal{X}_\ell$ for all k, and for each k there is a trajectory–action pair (y_k, a_k) such that $y_k(0) = x_k$, $y_k(t) \in \mathcal{X}_j$ for all $t \in [0, \eta]$, and

$$w(x_k) \leq \int_0^\eta u(y_k(t), a_k(t)) \exp(-\rho t)\,\mathrm{d}t + w(y_k(\eta)) \exp(-\rho\eta). \tag{29}$$

Moreover, from (25) we obtain that for k sufficiently large

$$v(x_k) \geq \int_0^\eta u(y_k(t), a_k(t)) \exp(-\rho t)\,\mathrm{d}t + v(y_k(\eta)) \exp(-\rho\eta). \tag{30}$$

Combining (29) and (30) yields

$$w(x_k) - v(x_k) \leq (w(y_k(\eta)) - v(y_k(\eta))) \exp(-\rho\eta) \leq M \exp(-\rho\eta).$$

Taking the limit $k \to \infty$ yields then $M \leq M \exp(-\rho\eta) < M$, again a contradiction. We conclude that necessarily $M \leq 0$. □

F.8 Proof of Proposition C.5

Proof. We give the proof for the subsolution case; the supersolution case is similar.

Set $\bar{x} = \bar{x}_j$. By hypothesis, the subsolution property holds for all $x \in \mathcal{X}_{j,+} \backslash \{\bar{x}\}$. Assuming that the statement of the proposition is false, there is a C^1 function ψ such that, firstly, $\psi(\bar{x}) = \bar{w}(\bar{x})$, secondly $\bar{w}(x) - \psi(x)$ restricted to $\mathcal{X}_{j,+}$ is maximal at $\bar{x}$, and finally

$$\rho\bar{w}(\bar{x}) - H_+(\bar{x}, \psi'(\bar{x})) > 0, \tag{31}$$

where $H_+ = H_{j+1}$. Introduce $\Delta(y) = \bar{w}(\bar{x} + y) - \psi(\bar{x} + y) - y^2$. Then Δ is continuous for $y \geq 0$, maximal at $y = 0$, and $\Delta(0) = 0$. Continuity implies that for every $n > 0$ there is $\xi_n > 0$ such that $\Delta(\xi_n) > -1/n$. On the other hand, if $y \geq 2/\sqrt{n}$, then $\Delta(y) \leq -y^2 \leq -4/n$. It follows that $0 < \xi_n < 2/\sqrt{n}$.

Set $\varepsilon_n = \xi_n/n$. The function

$$\Delta(y) - \varepsilon_n/y = \bar{w}(\bar{x}+y) - \big(\psi(\bar{x}+y) + y^2 + \varepsilon_n/y\big)$$

satisfies $\Delta(\xi_n) - \varepsilon_n/\xi_n \geq -2/n$ and $\Delta(y) - \varepsilon_n/y \leq -4/n$ if $y \geq 2/\sqrt{n}$. Hence it takes its maximum at a point $0 < y_n < 2/\sqrt{n}$, and, setting $x_n = \bar{x} + y_n$, we have

$$p_n \equiv \psi'(x_n) + 2y_n - \varepsilon_n/y_n^2 \in D^+\bar{w}(x_n).$$

As y_n maximises $\Delta(y) - \varepsilon/y$, we have, first, that $\Delta(y_n) - \varepsilon_n/y_n \geq \Delta(\xi_n) - \varepsilon_n/\xi_n \geq -2/n$, and, second, that $0 < \varepsilon_n/y_n \leq \Delta(y_n) + 2/n \leq 2/n$. Consequently, if $n \to \infty$ we obtain

$$p_n y_n = \psi'(x_n)y_n + 2y_n^2 - \varepsilon_n/y_n \to 0. \tag{32}$$

Since $v, w \in \mathscr{G}$ and one of these functions is discontinuous at $\bar{x}$, the point $\bar{x}$ is a right semi-repeller. In particular this implies that $(H_+)_p(\bar{x}, p) = f_+(\bar{x}, p) \geq 0$ for all p.

Writing $q_n = q^*(x_n, p_n)$, there are $0 < \theta_n^{(1)}, \theta_n^{(2)} < 1$ such that

$$\begin{aligned} H_+(x_n, p_n) &= u(\bar{x}, q_n) + u_x(\bar{x} + \theta_n^{(1)} y_n, q_n)y_n + p_n[f_+(\bar{x}, q_n) + (f_+)_x(\bar{x} + \theta_n^{(2)} y_n, q_n)y_n] \\ &= H_+(\bar{x}, p_n) + r_n \leq H_+(\bar{x}, \psi'(x_n) + 2y_n) + r_n, \end{aligned}$$

where we have set $r_n = u_x(\bar{x} + \theta_n^{(1)} y_n, q_n)y_n + (f_+)_x(\bar{x} + \theta_n^{(2)} y_n, q_n)p_n y_n$, and where we have used that $p_n \leq \psi'(x_n) + 2y_n$ as well as the fact that $H_+(\bar{x}, p)$ is nondecreasing in p. As u_x and $(f_+)_x$ are bounded, equation (32) implies that $r_n \to 0$ as $n \to \infty$.

Since w is a subsolution, we have $\rho w(x_n) \leq H_+(x_n, p_n) \leq H_+(\bar{x}, \psi'(x_n) + 2y_n) + r_n$. Taking the limit $n \to \infty$ then yields $\rho\bar{w}(\bar{x}) \leq H_+(\bar{x}, \psi'(\bar{x}))$, contradicting (31). $\square$

F.9 Proof of Proposition D.3

We need a technical result about linearisations (e.g. Cannarsa et al., 2015, Lemma 2.3).

Lemma F.2. *Let $g(t, x)$ be measurable in t and continuously differentiable in x. For $x \in \mathring{\mathfrak{X}}$, denote by $y(t; x)$ the solution to $\dot{y}(t) = g(t, y(t))$, $y(0) = x$. Assume that for $x_0 \in \mathring{\mathfrak{X}}$ there is $T > 0$ such that $y(t; x_0) \in \mathring{\mathfrak{X}}$ for all $t \in [0, T]$. Let Φ be the absolutely* *continuous solution of the linear system*

$$\dot{\Phi}(t) = g_x(t, y(t; x_0))\Phi(t), \quad \Phi(0) = 1.$$

Then for all x in a neighbourhood of x_0 in $\mathring{\mathfrak{X}}$, we have for $t \in [0, T]$

$$y(t; x) = y(t; x_0) + \Phi(t)(x - x_0) + o_t(|x - x_0|),$$

with $o_t(|x - x_0|)/|x - x_0| \to 0$ as $x \to x_0$, uniformly in t.

In the proof below, the *lower Dini directional derivative* is used, which for a continuous function $W(x)$ is defined as $\partial^- W(x; \xi) = \liminf_{h \downarrow 0}(W(x + h\xi) - W(x))/h$. Unlike an ordinary derivative, this derivative exists for all x and ξ. Clearly, if W is differentiable at x, then $\partial^- W(x; \xi) = W'(x)\xi$ for all ξ.

Proof of Proposition D.3. Let $I = y^*([0, T])$ be the orbit of the optimal trajectory. If y^* is constant, then I consists of a single point and V is differentiable on all of I. If y^* is non-constant, then I has positive length and by Proposition D.2 the value function V is differentiable on a dense subset $S \subset I$.

We first establish a relation between the derivatives V' on different points in S using a linearisation argument. Then we show that V' restricted to S is continuous, which will finally imply that V' exists everywhere in I.

For $z \in \mathring{\mathfrak{X}}$, let $y(t; z)$ and $\Phi(t)$ be, respectively, the solutions of $\dot{y}(t; z) = f(y(t; z), a^*(t))$ and $y(0; z) = z$, and of

$$\dot{\Phi}(t) = f_x(y(t; z), a^*(t))\Phi(t), \quad \Phi(0) = 1.$$

Choose $\xi \in \mathbb{R}$ arbitrarily, and take $h > 0$ such that $y(t; x + h\xi) \in \mathring{\mathfrak{X}}$ for all $t \in [0, T]$. By the optimality principle,

$$V(x + h\xi) \geq \int_0^t \exp(-\rho s) u(y(s; x + h\xi), a^*(s)) \,\mathrm{d}s + V(y(t; x + h\xi)) \exp(-\rho t).$$

For the optimal pair (y^*, c^*), we have

$$V(x) = \int_0^t \exp(-\rho s) u(y^*(s), a^*(s))\,\mathrm{d}s + V(y^*(t))\exp(-\rho t).$$

Differentiability of V at x implies

$$\begin{aligned}
V'(x)\xi &= \liminf_{h \downarrow 0} (V(x+h\xi) - V(x))/h \\
&\geq \liminf_{h\downarrow 0} \Bigg(h^{-1} \int_0^t \exp(-\rho s) \Big(u(y(s; x+h\xi), a^*(s)) - u(y^*(s), a^*(s)) \Big)\,\mathrm{d}s \\
&\qquad\qquad + \exp(-\rho t) \frac{V(y(t; x+h\xi)) - V(y^*(t))}{h} \Bigg) \\
&= \int_0^t \exp(-\rho s) u_x(y^*(s), a^*(s)) \Phi(s)\xi\,\mathrm{d}s + \exp(-\rho t) \partial^- V(y^*(t); \Phi(t)\xi),
\end{aligned}$$

where in the last equality Lemma F.2 has been used.

For $t \in [0, T]$ such that $y^*(t) \in S$, we find

$$V'(x)\xi \geq \int_0^t \exp(-\rho s) u_x(y^*(s), a^*(s)) \Phi(s)\xi\,\mathrm{d}s + \exp(-\rho t) V'(y^*(t)) \Phi(t)\xi.$$

Taking successively $\xi = 1$ and $\xi = -1$ yields

$$V'(x) = \int_0^t \exp(-\rho s) u_x(y^*(s), a^*(s)) \Phi(s)\,\mathrm{d}s + \exp(-\rho t) V'(y^*(t)) \Phi(t). \tag{33}$$

As $\Phi(t) \neq 0$, for all $t \in [0, T]$ we define a function $\hat{p}(t)$ by the relation

$$V'(x) = \int_0^t \exp(-\rho s) u_x(y^*(s), a^*(s)) \Phi(s)\,\mathrm{d}s + \exp(-\rho t) \hat{p}(t) \Phi(t).$$

Clearly $\hat{p}(0) = V'(x)$. Differentiating this relation with respect to t shows moreover that $\hat{p}$ satisfies (20), and consequently that $\hat{p}(t) = p^*(t)$ for all t. We then infer from (33) that $p^*(t) = V'(y^*(t))$ whenever $y^*(t) \in S$.

Take $z \in S$, and consider a sequence $z_n \in S$ such that $z_n \to z$ as $n \to \infty$. Find a sequence t_n such that $z_n = y^*(t_n)$. If necessary after passing to a subsequence we may assume—as $[0, T]$ is compact—that $t_n \to \tau$, and therefore $V'(z_n) = V'(y^*(t_n)) = p(t_n) \to p(\tau) = V'(z)$ as $n \to \infty$. Hence V' is continuous on S, and can uniquely be extended to a continuous function on I, which implies that V is continuously differentiable on I.

To prove the last statement, differentiate the equality

$$V(x) = \int_0^t \exp(-\rho s) u(y^*(s), a^*(s))\,\mathrm{d}s + \exp(-\rho t) V(y^*(t))$$

with respect to t and divide by $\exp(-\rho t)$ to obtain

$$\rho V(y^*(t)) = u(y^*(t), a^*(t)) + V'(y^*(t)) f(y^*(t), a^*(t)).$$

Since V is a supersolution and $p^*(t) = V'(y^*(t))$, we have $\rho V(y^*(t)) \geq H(y^*(t), p^*(t))$, which reads as

$$u(y^*(t), a^*(t)) + V'(y^*(t)) f(y^*(t), a^*(t)) \geq \max_q (u(y^*(t), q) + p^*(t) f(y^*(t), q)),$$

the last statement of the proposition. □

F.10 Proof of Proposition D.4

We begin by proving the first three statements of Proposition D.4.

Proof of Proposition D.4, Statements (i)–(iii).

(i) If y^* is not constant, then $\dot{y}^*(0) \neq 0$ and y^* is locally invertible on an interval $[0, \varepsilon_0)$. Consequently for every $0 < \varepsilon < \varepsilon_0$ there is $0 < t_1 < \varepsilon$ such that $y^*(t_1) \in \mathcal{D}_1$, as $\mathcal{D}_1$ is dense. But then $y^*(t) \in \mathcal{D}_1$ for all $t \geq t_1$ by Proposition D.3. As $\varepsilon > 0$ was arbitrary, this shows that $y^*(t) \in \mathcal{D}_1$ for all $t > 0$ such that $y^*(t) \in \mathring{\mathcal{X}}$.

(ii) Let $\tau > 0$ be such that $y^*(\tau) \in \mathcal{D}_1$. The trajectory (y^τ, p^τ) starting at the point $(y^*(\tau), V'(y^*(\tau)))$ satisfies $y^\tau(t) = y^*(\tau + t)$: in particular, we have $y^*(0) = y^\tau(-\tau)$ and $p_0 = p^\tau(-\tau)$.

(iii) Monotonicity of y^* has been shown in, e.g., Wagener (2003). □

In the proof of the last statement of Proposition D.4, and several other results, we shall use the invariant manifold theorem (Hirsch et al., 1977; Takens and Vanderbauwhede, 2010), which for a planar vector field ensures the existence of invariant curves that are tangent to the eigenspaces of a steady state.

More precisely, let $\bar{\zeta}$ be a steady state of a real analytic planar vector field $Y : \mathbb{R}^2 \to \mathbb{R}^2$, whose linearisation $DY(\bar{\zeta})$ has eigenvalues λ_1 and λ_2. If $\lambda_1 < 0 < \lambda_2$, there are unique and real analytic invariant manifolds W^{s} and W^{u} tangent, respectively, to the eigenspace E_1 of λ_1 and E_2 of λ_2 at the steady state, the *stable* and *unstable* manifolds. If $\lambda_1 = 0 < \lambda_2$, there exists a, not necessarily unique, C^∞ invariant manifold W^{cs} tangent to E_1 at the steady state, the *centre-stable* manifold, and a unique real analytic unstable manifold W^{u} tangent to E_2. If $0 < \lambda_1 < \lambda_2$, then there exists a unique and real analytic invariant manifold W^{uu} tangent to E_2 at the steady state, called the *strongly unstable* manifold. If $0 < \lambda_1 < \lambda_2$, all trajectories not on W^{uu} are tangent to E_1 and can be parametrised as the graph of a C^1 function $w : E_1 \to E_2$.

The eigenvalues and eigenspaces that correspond to a given invariant manifold are denoted by the same superscript: e.g. the centre-stable manifold W^{cs} is tangent to the centre-stable eigenspace E^{cs} of λ^{cs}.

We are mostly concerned with centre-stable manifolds. These manifolds are in general not unique and only infinitely often differentiable, not real analytic. However, the following result provides a condition for unicity and analyticity of the centre-manifold.

Theorem F.1 (Aulbach (1986))**.** *Let $\bar{\zeta}$ be a steady state of a real analytic vector field $Y : N \to \mathbb{R}^2$, where N is a neighbourhood of $\bar{\zeta}$ in $\mathbb{R}^2$. Let $\lambda_1 = 0$ and $\lambda_2 > 0$ be the eigenvalues of $DY(\bar{\zeta})$, and let E_1 and E_2 be the corresponding eigenspaces.*

If every neighbourhood of $\bar{\zeta}$ contains a fixed point of Y different from $\bar{\zeta}$, then there is a disk $D \subset N$ of positive radius, centred at $\bar{\zeta}$, and a unique analytic local centre–stable manifold $W^{\mathrm{cs}} \subset D$, tangent to E_1, such that all points on W^{cs} are steady states of Y.

The next result solves the Hamilton–Jacobi–Bellman equation if the action schedule takes a corner value q_b, $b \in \{\ell, u\}$, if we set $g(x) = f(x, q_b)$ and $v(x) = u(x, q_b)$.

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

$$\rho V(x) - v(x) - V'(x)g(x) = 0. \tag{34}$$

(i) Equation (34) *has bounded solutions V, all of which satisfy $V(\bar{x}) = v(\bar{x})/\rho$.*

(ii) If $g'(\bar{x}) < \rho$, each solution is continuously differentiable and $V'(\bar{x}) = v'(\bar{x})/(\rho - g'(\bar{x}))$.

(iii) If $g'(\bar{x}) < 0$, the solution V is unique and real analytic.

Proof. For g identically zero $V(x) = v(x)/\rho$ is the unique real analytic solution of (34).

Otherwise $\bar{x}$ is an isolated zero of g. Let $N = (x_1, x_2)$ be an open interval containing $\bar{x}$. Restricted to $N \times \mathbb{R}$, the graph of a differentiable solution V of (34) is a union of orbits of

$$\dot{y} = g(y), \quad \dot{w} = \rho w - v(y)$$

as $w(t) = V(y(t))$. This system has a unique steady state $(\bar{x}, \bar{w}) = (\bar{x}, v(\bar{x})/\rho)$. The linearisation $\begin{pmatrix} g'(\bar{x}) & 0 \\ -v'(\bar{x}) & \rho \end{pmatrix}$ at steady state has eigenvalues $\lambda_1 = g'(\bar{x})$ and $\lambda_2 = \rho > 0$, and corresponding eigenspaces $E_1 = \mathbb{R}\begin{pmatrix} \rho - g'(\bar{x}) \\ v'(\bar{x}) \end{pmatrix}$ and $E_2 = \mathbb{R}\begin{pmatrix} 0 \\ 1 \end{pmatrix}$.

If $\lambda_1 > 0$, the steady state is a repeller. For N sufficiently small there are two trajectories $(y_i(t), w_i(t))$, $i = 1, 2$, converging to the steady state as $t \to -\infty$, with $y_i(t_i) = x_i$ for some t_i. As $\dot{y}_i(t) \neq 0$ if $t < t_i$, these trajectories yield a continuous solution on N by setting $V(y_i(t)) = w_i(t)$ and $V(\bar{x}) = v(\bar{x})/\rho$. If $\lambda_1 \leq 0$, the centre-stable manifold W^{cs} of the steady state is tangent to E_1 and the graph of a bounded solution. This shows (i).

If $\lambda_1 < \rho$, the manifold E_2 is invariant, and any bounded solution trajectory not on E_2 is on a manifold W_1 tangent to E_1 at the steady state, which is the graph of a C^1 solution V to (34). The gradient of V at the steady state is the inclination of the eigenspace E_1, which evaluates to $V'(\bar{x}) = v'(\bar{x})/(\rho - g'(\bar{x}))$, showing (ii).

If $\lambda_1 < 0$, the manifold W^1 is the stable manifold of the steady state, which is unique and real analytic, completing the proof. □

Lemma F.3. *Let (y, p) be a canonical trajectory such that $y(t) \to \bar{x}_\pm$ as $t \to \pm\infty$. Then $p(t) \to \bar{p}_\pm$ as $t \to \pm\infty$, with $\bar{p}_\pm \in \mathbb{R} \cup \{-\infty, \infty\}$.*

Proof. We have to show that the set $\{t \geq 0 : p(t)\}$ has at most one accumulation point.

Assume that $\bar{p}_1 < \bar{p}_2$ are two distinct accumulation points. Then for all $n > 0$ there are $t_{1,n} < t_{2,n} < t_{1,n+1}$ such that $t_{1,n}$, $t_{2,n} \to \infty$ and $p(t_{1,n}) \to \bar{p}_1$, $p(t_{2,n}) \to \bar{p}_2$ as $n \to \infty$. Taking $p \in (\bar{p}_1, \bar{p}_2)$, for all $n > 0$ there are $t_{1,n} < t_{3,n} < t_{2,n} < t_{4,n} < t_{1,n+1}$ such that $p(t_{3,n}) = p(t_{4,n}) = p$ and $\dot{p}(t_{3,n}) > 0$ and $\dot{p}(t_{4,n}) < 0$. We conclude that the second component X_2 of the canonical vector field, and hence X itself, must vanish for all $(\bar{x}, p)$ with $p \in [\bar{p}_1, \bar{p}_2]$. As $\partial X_1/\partial p \neq 0$ on $\mathcal{P}_{\text{int}}$, we must have that $\bar{p}_1 \geq p_u(\bar{x})$ or $\bar{p}_2 \leq p_\ell(\bar{x})$. But for $(\bar{x}, p) \in \mathcal{P}_\ell \cup \mathcal{P}_u$, the conditions $X_2 = (\rho - f_x)p - u_x = 0$ and $\partial X_2/\partial p = (\rho - f_x) = 0$ imply $u_x = 0$, contradicting Assumption 2. The argument for $\{t \leq 0 : p(t)\}$ is analogous. □

Proof of Proposition D.4, Statement (iv). If the first alternative does not hold, we have $y^*(t) \in \mathcal{X}$ for all $t \geq 0$. Let (y, p) be the canonical trajectory such that $y^* = y$, whose existence is guaranteed by (ii). Then (iii) implies that $y(t) \to \bar{x} \in \mathcal{X}$ as $t \to \infty$.

Lemma F.3 implies that $p(t)$ converges to a limit as $t \to \infty$, or diverges to ∞, or to $-\infty$. Proposition A.3 rules out the second possibility. We prove the result by showing that the third possibility cannot occur either.

If $p(t) \to -\infty$ as $t \to \infty$, there is t_1 such that $(y(t), p(t)) \in \mathcal{P}_\ell$ for all $t \geq t_1$. Introduce $v(x) = u(x, q_\ell)$ and $g(x) = f(x, q_\ell)$. Then $H(x, p) = v(x) + pg(x)$ and $\dot{y}(t) = g(y(t))$ for all $t \geq t_1$. In particular $g(\bar{x}) = 0$ and $g'(\bar{x}) \leq 0$.

Set $I = y((t_1, \infty))$. By (i), on I the value function is differentiable and satisfies (34). Proposition F.1 implies that $V'(\bar{x}) = v'(\bar{x})/(\rho - g'(\bar{x})) \geq v'(\bar{x})/\rho$. As $p(t) \to V'(\bar{x})$ for $t \to \infty$, we have reached the desired contradiction. □

F.11 Proofs of Propositions D.5 and D.6

First we note an implication of real analyticity. For $b \in \{\ell, u\}$ and $z \in \mathcal{S}_b$, let $n_b(z)$ be the unit normal vector to $\mathcal{S}_b$ pointing out of $\mathcal{P}_{\text{int}}$ at z, and let $\omega_b(x) = n_b(x, p_b(x)) \cdot X(x, p_b(x))$.

Lemma F.4. *Let $b \in \{\ell, u\}$ and let $C \subset \mathring{\mathcal{X}}$ be a compact set.*

(i) If the cardinality of $\{x \in C : f(x, q_b) = 0\}$ is infinite, then $f(x, q_b) = 0$ for all $x \in \mathcal{X}$.

(ii) If the cardinality of $\{x \in C : \omega_b(x) = 0\}$ is infinite, then $\mathcal{S}_b$ is invariant under X.

Proof. Both assertions follow from the fact that a real analytic function whose zeros have an accumulation point vanishes identically. □

Proof of Proposition D.5. Assume the statement is false. Then for every $n > 0$ there are initial points $\tilde{x}_n$ located in different non-constant optimal orbits. We may assume that they are ordered in an increasing or decreasing sequence: we prove the result for the increasing case, the other being similar. By Proposition D.4.(ii) there are $\tilde{p}_{n,0}$ such that the non-constant trajectory starting at $\tilde{x}_n$ is the state component of the optimal canonical trajectory $(\tilde{y}_n, \tilde{p}_n)$ starting at $(\tilde{x}_n, \tilde{p}_{n,0})$. Proposition D.4.(iv) implies that this canonical trajectory converges to a point $(\bar{x}_n, \bar{p}_n)$ with $\bar{x}_n \in (\tilde{x}_n, \tilde{x}_{n+1})$; necessarily $\bar{p}_n \leq 0$.

Introduce $T_n = \min\{\tau \geq 0 : (\tilde{y}_n(t), \tilde{p}_n(t)) \in \mathcal{P}_b \text{ for all } t \geq \tau\}$. If $0 < T_n < \infty$ infinitely often, or after relabelling, for all $n > 0$, then the points $(\bar{x}_n, \bar{p}_n)$ are steady states of X in $\mathcal{P}_b$ and $f(\bar{x}_n, q_b) = 0$; applying Lemma F.4, it follows that $f(x, q_b) = 0$ for all x. But then $f_x(x, q_b) = 0$ for all x as well, and $X_2(x, p) = \rho p - u_x(x, q_b)$ if $(x, p) \in \mathcal{P}_b$. As $X_2(\bar{x}_n, \bar{p}_n) = 0$, it follows that $\omega_b(\bar{x}_n) < 0$. On the other hand, for $\hat{x}_n = \tilde{y}_n(T_n)$, we have $\omega_b(\hat{x}_n) \geq 0$. Hence ω_b vanishes in the interval $(\tilde{x}_n, \tilde{x}_{n+1})$; Lemma F.4 then implies that $\mathcal{S}_b$ is invariant. This however contradicts that $T_n > 0$.

If $T_n = 0$ infinitely often, it follows as above that $f(x, q_b) = 0$ for all x and $\tilde{y}_n(t) = \tilde{x}_n$ for all t, contradicting that $\tilde{x}_n$ is located in a non-constant optimal orbit.

Consider next the situation that $T_n = \infty$ for all n sufficiently large and $(\bar{x}_n, \bar{p}_n) \notin \mathcal{P}_{\text{int}}$. This implies that $(\bar{x}_n, \bar{p}_n) \in \mathcal{S}_b$. Since these points are steady states, $\omega_b(\bar{x}_n) = 0$, and by Lemma F.4 the set $\mathcal{S}_b$ is invariant and $f(x, q_b) = 0$ for all x. But then $\mathcal{S}_b$ is the centre–stable manifold for $(\bar{x}_n, \bar{p}_n)$, and there is no non-constant trajectory that tends to $(\bar{x}_n, \bar{p}_n)$ as $t \to \infty$, which contradicts the choice of $(\tilde{y}_n, \tilde{p}_n)$.

We are left with the situation that $(\bar{x}_n, \bar{p}_n) \in \mathcal{P}_{\text{int}}$ for n sufficiently large. A subsequence of these points converges to a steady state $(\bar{x}, \bar{p}) \in \overline{\mathcal{P}_{\text{int}}}$ of X. Hence by the Aulbach theorem, the invariant centre-stable manifold of $(\bar{x}, \bar{p})$ consists of steady states and contains $(\bar{x}_n, \bar{p}_n)$ for n sufficiently large, again contradicting the fact that $(\tilde{y}_n, \tilde{p}_n)$ is non-constant. □

Proposition F.2. *A compact set $C \subset \mathring{\mathcal{X}}$ contains only finitely many switching points of any optimal orbit.*

Proof. Let I be a non-constant optimal orbit, $b \in \{\ell, u\}$, $y : \mathcal{T} \to I$ a state trajectory parametrising I, and p such that (y, p) is the canonical trajectory associated to y.

Assume that there are infinitely many switching points in I. There is an increasing sequence $t_1 < t_2 < \ldots$ in $\mathcal{T}$ such that $(y(t_{2k-1}), p(t_{2k-1})) \in \mathcal{P}_{\text{int}}$ and $(y(t_{2k}), p(t_{2k})) \in \mathcal{P}_b$ for all $k > 0$. Consequently, there are times $t_1 < \hat{t}_1 < t_2 < \hat{t}_2 < \ldots$ such that

$(y(\hat{t}_n), p(\hat{t}_n)) = (\hat{x}_n, p_b(\hat{x}_n)) \in \mathcal{S}_b$, and such that $\omega_b(\hat{x}_{2k-1}) \geq 0$ and $\omega_b(\hat{x}_{2k}) \leq 0$ for all k. Hence ω_b vanishes in the interval $(\hat{x}_n, \hat{x}_{n+1})$ for every $n \geq 0$. By Lemma F.4, the set S_b is invariant, which contradicts the existence of switching points. □

Proof of Proposition D.6. If $x_0 \in \mathring{I} \cap \mathring{\mathcal{X}}_j$, then there is a non-constant optimal trajectory $y : \mathcal{T} \to I$, with $\mathcal{T} = \mathbb{R}$ or $\mathcal{T} = [0, \infty)$, such that $x_0 = y(t_0)$ with $t_0 \in \mathring{\mathcal{T}}$. By Proposition D.4, V is differentiable in a neighbourhood of x_0 and there is a canonical trajectory (y, p) such that $V'(y(t)) = p(t)$ for all $t > 0$ and such that $\dot{y}(t)$ does not change sign. Moreover, if x_0 is not a switching point, then $(y(t), p(t))$ is real analytic for t close to t_0, since it is locally the trajectory of a real analytic vector field X. Hence we can solve $x = y(t)$ as $t = y^{-1}(x)$ around x_0, and obtain that $V'(x) = p(t) = p(y^{-1}(x))$ is real analytic.

It remains to show that ϕ can be extended to a differentiable function on an open interval containing I. Let $\bar{t} \in \partial\mathcal{T}$: that is, $\bar{t} \in \{0, -\infty, \infty\}$. Proposition D.4.(ii) and Lemma F.3 imply that $(y(t), p(t))$ converges either to $(\bar{x}, \infty)$, or to $(\bar{x}, -\infty)$, or to a finite limit $(\bar{x}, \bar{p})$, as $t \to \bar{t}$. In the first and second case we respectively have $\phi(y(t)) = q^*(y(t), p(t))$ $q^*(y(t), p(t)) = q_u$ and $q^*(y(t), p(t)) = q_\ell$ for t in a neighbourhood of $\bar{t}$, and it is clear that ϕ can be differentiably extended.

In the third case $(y(t), p(t))$ tends to a steady state $\bar{z} = (\bar{x}, \bar{p})$ of the canonical vector field as $t \to \bar{t}$. Lemma F.2 implies that there is $0 < t_1 < \bar{t}$ such that $z(t) = (y(t), p(t))$ does not pass through a switching point for $t \in (t_1, \bar{t})$. If $z(t) \in \mathcal{P}_b$ for $b \in \{\ell, u\}$ for $t \in N$, then $\phi(y(t)) = q_b$ for those values of t, and we conclude as above. If $z(t) \in \mathcal{P}_{\text{int}}$ for all $t \in N$, the trajectory is tangent to an eigenspace E of $DX_I(\bar{z})$, where X_I is the real analytic extension to $\mathcal{X} \times \mathbb{R}$ of the restriction of X to $\mathcal{P}_{\text{int}}$. The fact that $H_{pp}(\bar{z}) > 0$ implies, first, that none of these eigenspaces is vertical, and, second, if $DX_I(\bar{z})$ has an eigenvalue with algebraic multiplicity 2, then the geometric multiplicity is 1.

As in all cases the eigenspaces are one-dimensional non-vertical lines, it follows that $V''(y(t)) = X_2(z(t))/X_1(z(t))$ converges to the inclination $\bar{w}$ of E with respect to the horizontal axis as $t \to \bar{t}$. Consequently $\phi'(y(t)) = q_x^*(y(t), p(t)) + q_p^*(y(t), p(t))V''(y(t))$ converges to $q_x^*(\bar{z}) + q_p^*(\bar{z})\bar{w}$. □

F.12 Proof of Proposition E.1

Proof. The dynamics take the form $\dot{y}(t) = \mathbf{f}^{\phi}(y(t))$. Given a trajectory y with initial state x, the payoff at x is

$$V_i^{\phi}(x) = \int_0^{\Theta} \exp(-\rho t)\mathbf{u}^{\phi}(y(t))\,\mathrm{d}t + \exp(-\rho\Theta)\beta(y(\Theta)).$$

Properties (i), (ii) and (iii) are immediate.

As f_j^{ϕ} is real analytic on $\mathring{\mathcal{X}}_j$, we either have that $f_j^{\phi}(x) = 0$ for all $x \in \mathcal{X}_j$, or the set $\mathcal{E}_j = \left\{x \in \mathcal{X}_j : f_j^{\phi}(x) = 0\right\}$ is finite. We set $\mathcal{E} = \cup_j \mathcal{E}_j$. If f_j^{ϕ} is identically zero, then $V_i^{\phi}(x) = u_{i,j}^{\phi}(x)/\rho$ is real analytic on $\mathcal{X}_j$. If not, take $x \in \mathcal{X}_j \backslash \mathcal{E}_j$. As $f_j^{\phi}(x) \neq 0$, we have

$$\begin{aligned} V_i^{\phi}(x + f_j^{\phi}(x)t + o(t)) - V_i^{\phi}(x) &= V_i^{\phi}(y(t)) - V_i^{\phi}(x) \\ &= (\exp(\rho t) - 1)V_i^{\phi}(x) - \exp(\rho t)\int_0^t \exp(-\rho s)\mathbf{u}_i^{\phi}(y(s))\,\mathrm{d}s, \end{aligned}$$

which on dividing by t and taking the limit $t \to 0$ yields, first, that the limit of the left hand expression exists, and, second, that it equals

$$\left(V_i^{\phi}\right)'(x)\, f_j^{\phi}(x) = \rho V_i^{\phi}(x) - u_{i,j}^{\phi}(x). \tag{35}$$

Note that the graph of solutions of (35) consists of trajectories of the dynamical system

$$\dot{y} = f_j^{\phi}(y), \quad \dot{v} = -u_{i,j}^{\phi}(y) + \rho v. \tag{36}$$

Then Proposition F.1 implies that V_i^{ϕ} is continuous on $\mathring{\mathcal{X}}_j$ and real analytic on $x \in \mathring{\mathcal{X}}_j \backslash \mathcal{E}$, showing (iv). Now (v) is also straightforward.

Let $\bar{x}$ be such that $f_j^{\phi}(\bar{x}) = 0$. If $(f_j^{\phi})'(\bar{x}) < \rho$, V_i^{ϕ} is differentiable at $\bar{x}$ by Proposition F.1. If $(f_j^{\phi})'(\bar{x}) \geq \rho$, trajectories of (36) tending to $(\bar{x}, u_{i,j}^{\phi}(\bar{x})/\rho)$ are tangent to an eigenspace of the linearisation of (36) at the steady state, showing that the limit of $(V_i^{\phi})'(x)$ as $x \to \bar{x}$ exist, even if it is possibly infinite. This shows (vi). $\square$