

Continuation Geometry

Flyxion

July 2026

Contents

1	Why Continuation Matters	2
2	Precursors	5
3	Viability Theory	10
4	The Selection Already Inside Viability	15
5	Continuation Operators	19
6	Continuation Spaces	23
7	Distinguishability and Repair	27
8	Admissibility Functionals	30
9	Preference Fields and Coordination Geometry	34
10	Small Machines	38
11	Closing Sketches, and What Was Earned	42
	Appendices	45
A	Support Functions and Lipschitz Estimates	46
B	Measurable Selections	47
C	Grönwall-Type A Priori Bounds	48
	References	49

Chapter 1

Why Continuation Matters

A system that has stopped changing has not necessarily stopped existing, but a system that has run out of ways to change has. There is a difference between a room that is quiet because nothing needs to happen in it and a room that is quiet because nothing can. The first is rest. The second is the end of a story that may not announce itself as an ending. Most of what concerns us about a living body, a piece of software, an institution, or a plan is not its condition at the present instant but the width of what it can still become. A body in a coma and a body in deep sleep can look alike from across the room; what distinguishes them is not visible in a single frame but in the shape of what still lies ahead of each.

This is a book about that shape. What a system could still do matters more than what it is doing right now: the whole reachable fan of futures opening out from wherever it happens to stand. Call that fan a system's futures, or its options, or its story-space. Whatever you call it, it's the thing that shrinks when a hinge rusts shut, when a species loses a gene it won't get back, when an argument commits itself to a position it can't walk back from, when a bridge corrodes past the point where any repair schedule could save it. Repair, in the most general sense of the word, tries to restore it: not a particular outcome, but the standing possibility of many outcomes.

Several habits of thought already study pieces of this. Engineers ask whether a structure will hold under load. Ecologists ask whether a population can recover from a disturbance. Economists ask whether an allocation can be sustained, and control theorists ask whether a system can be steered to stay within safe bounds indefinitely. Each of these is really the same question in different clothes: whether the fan of futures remains nonempty, whether there is still at least one way forward that respects whatever constraints the situation imposes. That question has a mature and exact mathematical answer, developed over several decades under the name of viability theory, and this book leans on that answer throughout.

But existence is not the only thing worth asking about a fan of futures. Two systems can both have somewhere left to go and still differ enormously in how much room they have to maneuver, how easily they could be nudged back from the edge, which of their internal distinctions survive into their futures and which quietly get erased, and which of their many possible continuations is the one actually worth choosing. A patient who can still recover and a patient who can barely still recover are both, technically, viable; no doctor would treat that technicality as the end of the conversation. The questions that begin where mere existence leaves off, how much room, how easily recovered, what is preserved, what is worth choosing, are what this book is organized around. They do not yet have as settled a mathematical treat-

ment as the existence question does. Building one, on top of the existence question rather than instead of it, is the project.

The order of concerns runs roughly as follows, and will recur throughout the book. A system draws distinctions: it treats some states as meaningfully different from others. Those distinctions are carried by information, and information degrades. This is what entropy names. Some of that degradation can be undone; the undoing is repair. Repair is only possible where a future still exists to be repaired into; that existence is continuation. And among the continuations that exist, some are better than others by standards that have nothing to do with mere existence; sorting them out is admissibility. Existence, in other words, is necessary for everything downstream of it but is far from sufficient for any of it. This book starts at existence, because it is the best-understood link in the chain, and spends the rest of its length building outward from it toward repair, distinction, and admissibility: the concerns that motivated picking up the chain in the first place.

The object of study

We now fix, once and for all, the formal object the rest of the book is about.

Definition 1.1 (State, dynamics, constraint). *Let X be a state space (a metric space; in applications typically \mathbb{R}^n or a finite-dimensional manifold). Let $F : X \rightrightarrows X$ be a set-valued map assigning to each state $x \in X$ a set $F(x) \subseteq X$ of admissible velocities, so that trajectories of the system are absolutely continuous curves γ satisfying the differential inclusion $\dot{\gamma}(t) \in F(\gamma(t))$. Let $K \subseteq X$ be closed; K is the constraint set, the region within which the system is required to remain.*

Definition 1.2 (Continuation space). *For $x \in K$, the continuation space of x is*

$$\Gamma_K(x) = \{\gamma : [0, \infty) \rightarrow K \mid \gamma(0) = x, \dot{\gamma}(t) \in F(\gamma(t)) \text{ for a.e. } t \geq 0\}.$$

For a finite horizon $T > 0$, the finite-horizon continuation space is

$$\Gamma_K^T(x) = \{\gamma : [0, T] \rightarrow K \mid \gamma(0) = x, \dot{\gamma}(t) \in F(\gamma(t)) \text{ for a.e. } t \in [0, T]\}.$$

$\Gamma_K(x)$ (or simply $\Gamma(x)$ where K is fixed by context) is the fan of futures referred to above, made precise relative to one particular reading of “possible”: a future is possible if it is a trajectory of the given dynamics that never leaves the given constraint set. Other readings of possible are available and will matter later: controllably possible, admissibly possible, epistemically possible. But this dynamical reading is where the book begins, because it is the one with an existing, exact theory attached to it.

Proposition 1.1 (Restriction is well-defined). *For every $T > 0$, restriction of a trajectory to $[0, T]$ defines a map*

$$\rho_T : \Gamma_K(x) \rightarrow \Gamma_K^T(x), \quad \rho_T(\gamma) = \gamma|_{[0, T]}.$$

Proof. If $\gamma \in \Gamma_K(x)$ then $\gamma(0) = x$, $\gamma(t) \in K$ for all $t \geq 0$, and $\dot{\gamma}(t) \in F(\gamma(t))$ for a.e. $t \geq 0$. Restricting to $[0, T]$ preserves each of these three conditions on the smaller domain, so $\gamma|_{[0, T]} \in \Gamma_K^T(x)$. \square

Proposition 1.2 (Consistency across horizons). *For $0 < T_1 < T_2$, restriction further to $[0, T_1]$ defines a map $\rho_{T_1, T_2} : \Gamma_K^{T_2}(x) \rightarrow \Gamma_K^{T_1}(x)$, and the diagram of restriction maps commutes: $\rho_{T_1} = \rho_{T_1, T_2} \circ \rho_{T_2}$ for every $T_1 < T_2 < \infty$.*

Proof. Immediate from the definition of restriction: restricting a curve to $[0, T_1]$ gives the same result whether it is done directly from the domain $[0, \infty)$ or via an intermediate restriction to $[0, T_2]$, since restriction of a function to a subset of its domain depends only on that subset, not on the route by which the subset was reached. \square

Remark 1.1. *This consistency is what allows $\Gamma_K(x)$ to be recovered, in good cases, as the projective (inverse) limit of the finite-horizon spaces $\Gamma_K^T(x)$ as $T \rightarrow \infty$. This observation is recorded here because Chapter 6 will build continuation volume and continuation distance on the finite-horizon spaces precisely because they are well-posed where the infinite-horizon trajectory space is not; the present proposition is what justifies treating the finite-horizon construction as an approximation to $\Gamma_K(x)$ rather than a different object entirely.*

Proposition 1.3 (Equilibria have nonempty continuation spaces). *If $x^* \in K$ satisfies $0 \in F(x^*)$, then $\Gamma_K(x^*) \neq \emptyset$.*

Proof. Let $\gamma(t) = x^*$ for all $t \geq 0$. Then $\gamma(0) = x^*$, $\gamma(t) = x^* \in K$ for all t , and $\dot{\gamma}(t) = 0 \in F(x^*) = F(\gamma(t))$ for all t . Hence $\gamma \in \Gamma_K(x^*)$. \square

Remark 1.2. *This is the first, and simplest, example of a nonempty continuation space, and it is worth pausing on because it already illustrates a distinction the rest of the book will not let go of: nonemptiness of $\Gamma_K(x^*)$ here is witnessed by a single, unmoving trajectory. It says nothing about how many other continuations x^* has, how close x^* sits to states with no continuations at all, or whether the distinctions the system cares about survive along the witnessing trajectory. A full characterization of nonemptiness for arbitrary $x \in K$, not just equilibria, is the content of Aubin's Viability Theorem, taken up next.*

The progression this book follows

The chain of concerns described in natural language above, distinction, information, entropy, repair, continuation, admissibility, is not a theorem and is not offered as one. It is a roadmap, and its purpose is to say in advance which part of the whole problem each part of the book is responsible for.

Continuation, in the sense of $\Gamma_K(x) \neq \emptyset$, is the best-understood link and is addressed first: it is the subject of viability theory, surveyed and put to use in the chapters immediately following this one. Repair is addressed once continuation is in hand, since a repair is, formally, nothing more than a continuation that reaches a designated target: the machinery for saying this precisely is capture, and it costs nothing beyond what continuation has already supplied. Distinction and admissibility are addressed last, and are the hardest part of the roadmap to formalize, because they require deciding not merely whether a future exists but which of the many existing futures should be preferred, and by what standard. The book's later chapters build that standard in stages rather than assuming it in advance.

Chapter 2

Precursors

No idea arrives without ancestors, and an idea that pretends otherwise usually turns out to be a poorer version of something already on the shelf. The temptation for any project that finally finds its central object is to treat the discovery as an arrival rather than a continuation of its own, to write the history chapter last, briefly, and mostly as a courtesy. That temptation is worth resisting here for a specific reason: the people who spent real years on the questions this book asks were not working on one shared problem that only now gets its proper name. They were working on at least two different problems that happen, from a certain distance, to look like the same problem.

The first tradition asks how a system manages to keep going. Given where a system stands and how it is allowed to move, will it still be able to move tomorrow, and the day after, without breaking whatever rule keeps it a functioning version of itself? This is a question about persistence, and it has been asked with increasing precision for over a century: orbits that never quite settle down, equilibria that recover from being nudged, decisions that must be made in sequence without knowing the whole future in advance. Each advance in this tradition sharpened the same underlying question rather than replacing it, which is why it reads, in retrospect, like a single conversation carried on by different people.

The second tradition asks a different question, one that does not reduce to the first even though it sounds similar. Given where a system stands, what does the whole space of things it could still become look like? Not whether that space is empty, but what shape it has, how it grows or shrinks as choices get made, and which of its features are worth preserving. This tradition is younger, more scattered across disciplines that rarely cite one another, and considerably less tidy. Evolving biological form. What innovations become available once other innovations have already happened. Causal alternatives to what actually occurred. Even an architect's attempt to say which building layouts leave a neighborhood room to keep being a good place to live. None of these projects thought of itself as contributing to a general theory of possibility-spaces. Read together, they look like exactly that.

This book sits where these two traditions meet, and owes each of them a debt it should not obscure. What follows makes some of that debt precise, showing in each case how much of the earlier work can be recovered as a special case or a direct ancestor of the object this book is built around, and where each earlier tradition stops short.

The persistence lineage

Definition 2.1 (Recurrent state, after Poincaré). Let $\phi_t : X \rightarrow X$ be a measure-preserving flow on a probability space (X, μ) , and let $A \subseteq X$ have $\mu(A) > 0$. A point $x \in A$ is recurrent in A if $\phi_t(x) \in A$ for some $t > 0$, and infinitely often for a.e. such x .

Proposition 2.1 (Poincaré recurrence and nonempty continuation). If A is invariant under ϕ_t in the sense that $\phi_t^{-1}(A)$ has full relative measure in A for all t , then for a.e. $x \in A$, the trajectory $\gamma(t) = \phi_t(x)$ lies in $\Gamma_A(x)$, and moreover returns to A infinitely often.

Proof. This is the Poincaré Recurrence Theorem, applied to the flow restricted to A : almost every point of a measure-preserving invariant set returns to that set infinitely often under forward iteration [17]. The trajectory $\gamma(t) = \phi_t(x)$ satisfies $\gamma(0) = x$ and $\gamma(t) \in A$ for all $t \geq 0$ by invariance of A , and $\dot{\gamma}(t)$ equals the generator of ϕ_t at $\gamma(t)$, which lies in $F(\gamma(t))$ trivially when F is taken to be the vector field generating ϕ_t . Hence $\gamma \in \Gamma_A(x)$. \square

Remark 2.1. This is the earliest instance, in the lineage this book claims, of a proof that a continuation space is nonempty, though Poincaré's concern was not persistence under constraint but the much stronger fact of eventual return. Recurrence is strictly stronger than mere viability: every recurrent trajectory witnesses nonemptiness of $\Gamma_A(x)$, but nonemptiness of $\Gamma_K(x)$ does not require any trajectory to return to its starting neighborhood at all. The relation between recurrence and viability, when a viable trajectory can be strengthened to a recurrent one, is not pursued in this book but is flagged here as a legitimate question left on the table by folding Poincaré into the same lineage as Aubin.

Definition 2.2 (Lyapunov stability). A state x^* with $0 \in F(x^*)$ is Lyapunov stable if for every $\varepsilon > 0$ there exists $\eta > 0$ such that $\|x - x^*\| < \eta$ implies every $\gamma \in \Gamma_X(x)$ satisfies $\|\gamma(t) - x^*\| < \varepsilon$ for all $t \geq 0$.

Proposition 2.2 (Lyapunov functions certify viability of sublevel sets). Suppose $V : X \rightarrow \mathbb{R}_{\geq 0}$ is continuously differentiable, $V(x^*) = 0$, $V(x) > 0$ for $x \neq x^*$, and $\sup_{v \in F(x)} \nabla V(x) \cdot v \leq 0$ for all x . Then for every $c > 0$, the sublevel set $K_c = \{x : V(x) \leq c\}$ is viable under F : $\Gamma_{K_c}(x) \neq \emptyset$ for every $x \in K_c$.

Proof. Fix $x \in K_c$ and let γ solve $\dot{\gamma}(t) \in F(\gamma(t))$, $\gamma(0) = x$ (such a solution exists locally under standard regularity hypotheses on F). Along γ , $\frac{d}{dt} V(\gamma(t)) = \nabla V(\gamma(t)) \cdot \dot{\gamma}(t) \leq \sup_{v \in F(\gamma(t))} \nabla V(\gamma(t)) \cdot v \leq 0$, so $V(\gamma(t))$ is nonincreasing. Since $V(\gamma(0)) = V(x) \leq c$, it follows that $V(\gamma(t)) \leq c$ for all t , i.e. $\gamma(t) \in K_c$ for all t . Hence $\gamma \in \Gamma_{K_c}(x)$. \square

Remark 2.2. This is the classical Lyapunov stability argument [14], restated to make explicit what it already was: a direct construction of a nonempty, in fact tangentially controlled, continuation space. The tangency criterion of Aubin's Viability Theorem can be read as the generalization of this proposition obtained by dropping the requirement that a single global function V certify viability for every sublevel set at once, and instead checking the weaker, purely local condition $F(x) \cap T_K(x) \neq \emptyset$ directly on the boundary of whatever set K is given, without needing K to arise as a sublevel set of anything.

Definition 2.3 (Bellman's principle of optimality). For a controlled system with running cost l and terminal cost c , define the value function

$$J(t, x) = \inf_{u(\cdot)} \left[\int_t^T l(x(s), u(s)) ds + c(x(T)) \right]$$

over admissible controls $u(\cdot)$ steering $\dot{x}(s) = f(x(s), u(s))$ from $x(t) = x$.

Proposition 2.3 (Dynamic programming as a special case of guaranteed valuation). *Suppose $Q(x) = \{0\}$ (no perturbation) in Aubin’s dynamic core construction. Then Aubin’s guaranteed valuation function $V^\sharp(T, x)$ reduces to Bellman’s value function $J(t, x)$ under the correspondence $T \leftrightarrow T - t$, and the Hamilton-Jacobi-Isaacs variational inequality reduces to the Hamilton-Jacobi-Bellman equation*

$$-\frac{\partial J}{\partial t} + \inf_{u \in P(x)} [\nabla J(t, x) \cdot f(x, u) + l(x, u)] = 0.$$

Proof. When $Q(x) = \{0\}$, the supremum over perturbations $v \in Q(x)$ in the defining formula for V^\sharp (Aubin, Theorem 1.2) has a single term to range over, so it drops out, leaving only the infimum over feedbacks $p_e(x) \in P(x)$. This is exactly Bellman’s infimum over controls, and the discount term $m(x, p, v)$, in the absence of perturbation, plays the role of a running cost once l is identified with the appropriate component of the functional J_u . Substituting into the Hamilton-Jacobi-Isaacs inequality and dropping the now-vacuous $\sup_{v \in Q(x)}$ term gives the stated Hamilton-Jacobi-Bellman equation [6]. \square

Remark 2.3. *Dynamic programming is, on this reading, the special case of guaranteed valuation in which nature is not permitted to intervene. Aubin’s contribution to the persistence lineage was to reinstate the adversary that Bellman’s formulation sets to zero, without discarding the optimization structure Bellman had already built. Aubin’s viability kernel and capture basin are the further generalization obtained by dropping the requirement that a single scalar cost be optimized at all, and asking only whether some admissible control keeps the system where it needs to be: existence in place of optimality, which is where this book’s own starting point sits.*

The possibility-space lineage

Definition 2.4 (Adjacent possible, after Kauffman). *Given a set of currently realized configurations $S_t \subseteq X$ and a one-step generative map $\Phi : X \rightarrow 2^X$ (the set of configurations directly reachable from a given one), the adjacent possible of S_t is*

$$\text{Adj}(S_t) = \bigcup_{x \in S_t} \Phi(x).$$

Proposition 2.4 (Adjacent-possible growth as reachable-set recursion). *Define $S_{t+1} = S_t \cup \text{Adj}(S_t)$. Then, taking Φ as a discrete-time one-step reachability map and t discrete steps,*

$$S_t = \bigcup_{x \in S_0} R_t(x),$$

where $R_t(x)$ is the discrete-time finite-horizon reachable set generated by Φ from x .

Proof. By induction on t . The base case $t = 0$ is immediate: $R_0(x) = \{x\}$ and $S_0 = \bigcup_{x \in S_0} \{x\}$. For the inductive step, assume $S_t = \bigcup_{x \in S_0} R_t(x)$. Then

$$S_{t+1} = S_t \cup \text{Adj}(S_t) = \bigcup_{x \in S_0} R_t(x) \cup \bigcup_{y \in S_t} \Phi(y) = \bigcup_{x \in S_0} \left(R_t(x) \cup \bigcup_{y \in R_t(x)} \Phi(y) \right) = \bigcup_{x \in S_0} R_{t+1}(x),$$

using that $\Phi(y)$ for $y \in R_t(x)$ is exactly the one-step extension of a t -step reachable trajectory from x , which is what $R_{t+1}(x)$ collects. \square

Remark 2.4. *Kauffman’s adjacent possible [11] is, under this correspondence, an early and purely combinatorial instance of the finite-horizon reachable-set construction that later chapters use to build continuation volume. What it lacks, by design, is any constraint set K : the adjacent possible is not asked to remain within bounds, only to grow, so it corresponds to the unconstrained case $K = X$. It is this unconstrained framing, uncommon in the persistence lineage but common in Kauffman’s and Holland’s own writing [10], that marks the possibility-space lineage as a genuinely separate tradition rather than a duplicate of the first.*

Definition 2.5 (Structural causal model and intervention, after Pearl). *A structural causal model is a tuple (V, U, \mathcal{F}) of endogenous variables V , exogenous variables U , and structural equations \mathcal{F} , each $f_i \in \mathcal{F}$ determining V_i as a function of its parents and relevant exogenous noise. An intervention $\text{do}(V_j = v)$ replaces f_j with the constant function v , producing a modified model $(V, U, \mathcal{F}_{\text{do}(V_j=v)})$.*

Proposition 2.5 (Counterfactual reachability as continuation under intervened dynamics). *Let F correspond to the structural equations \mathcal{F} in the sense that $F(x)$ encodes the admissible next-states consistent with \mathcal{F} from configuration x , and let $F_{\text{do}(V_j=v)}$ be the corresponding map under the intervened model. Then the counterfactual “what would have happened to x had V_j been set to v ” corresponds exactly to $\Gamma_K^{F_{\text{do}(V_j=v)}}(x)$, the continuation space of x computed under the intervened dynamics rather than the original one.*

Proof. By construction: a counterfactual trajectory in Pearl’s sense is obtained by holding the exogenous noise U fixed at its actual values and re-evaluating the structural equations under the modified $\mathcal{F}_{\text{do}(V_j=v)}$ [16]. This is precisely the definition of a trajectory of the continuation space $\Gamma_K^{F_{\text{do}(V_j=v)}}(x)$: an admissible evolution from x under the dynamics $F_{\text{do}(V_j=v)}$ rather than F . \square

Remark 2.5. *This is the cleanest correspondence in the possibility-space lineage, and it points at something later chapters will need directly: comparing continuations across different choices of dynamics F , not merely across different starting states x under one fixed F . Every construction introduced so far has been stated for a single fixed (F, K) . Pearl’s framework is a reminder that the more general comparison, across interventions on F itself, is both meaningful and already partly worked out elsewhere; this book does not take it up in full, but the treatment of multiple agents sharing a dynamics later on is the nearest point of contact.*

Remark 2.6 (Levin: morphological and bioelectric continuation). *Levin’s work on bioelectric signaling as a substrate for the storage and recall of anatomical target-morphologies [13] is included in this lineage on different grounds than the constructions above: not because it yields a theorem restatable in the $\Gamma_K(x)$ vocabulary without loss, but because it is the clearest existing example of a system whose continuation space genuinely depends on more than its instantaneous state: a regenerating organism’s available futures depend on bioelectric pattern memory that persists through and beyond local tissue damage, in a way not recoverable from cell-by-cell instantaneous configuration alone. This is, in the vocabulary already fixed, exactly the history-dependence flagged as an open problem for the continuation-operator construction. No proposition is offered here forcing Levin’s biology into the present formalism; the honest claim is narrower, that this is the strongest existing example motivating why history-dependence needs to be taken seriously rather than assumed away.*

Remark 2.7 (Alexander: pattern languages). *Christopher Alexander’s pattern language [5] is the loosest member of this lineage and is kept loose deliberately. A pattern, in Alexander’s sense, is a reusable*

solution to a recurring design problem, stated together with the conditions under which it applies and the further patterns it tends to invite or foreclose. Read architecturally rather than dynamically, a choice of pattern is a move that changes which further patterns remain available, which is a possibility-space claim in substance, even though Alexander never formalized it as one and this book will not retroactively supply the formalization on his behalf. What is worth taking from Alexander is not a theorem but a discipline: that a possibility-space can be studied through the patterns that compose it and the compatibility relations between them, prior to and independent of any measure, metric, or dynamics being placed on it. That discipline anticipates the lattice-based treatment of admissibility taken up later, where composability and compatibility, rather than a single scalar score, are again the primary structure.

Where the lineages meet

The two lineages traced here answer different questions and were not, in general, in conversation with one another. The persistence lineage answers whether a future exists; its instruments, recurrence, stability, dynamic programming, viability, get sharper across a century of work without changing what they are fundamentally asking. The possibility-space lineage answers what the space of futures looks like once one exists; its instruments are more heterogeneous because the question is less settled, and in at least two cases above (Levin, Alexander) the correspondence to the present formalism is offered as motivation rather than as an established equivalence. The chapters that follow begin by finishing what the persistence lineage started: stating Aubin's Viability Theorem properly and using it as the foundation everything else stands on, before turning to the harder task of giving the possibility-space lineage's questions the same kind of precision the persistence lineage has already earned.

Chapter 3

Viability Theory

Walk along the top of a sea wall and there is a moment, well before you lose your footing, when the only directions left that keep you on the wall are a narrower and narrower band around straight ahead. You do not fall the instant that band closes to nothing; you fall the instant after, when a gust or a stumble sends you in some direction the band no longer contains. Whether you are in danger, in other words, is not really a question about your current position. Position alone cannot distinguish the middle of a wide, flat wall from its crumbling edge if both happen to put your feet at the same height. What matters is narrower than position: it is whether the directions physics is about to hand you and the directions that keep you on the wall still overlap.

This is a more useful fact than it first appears, because it turns a question about the entire rest of your walk (will I, at any point in the future, stumble off this wall) into a question you can answer by looking only at where your feet are right now and which way you could step next. The whole of the future collapses into a local, immediate test. That collapse is not a metaphor; it is a theorem, discovered first for smooth boundaries by Nagumo in the 1940s and given its general, set-valued form by Aubin and Frankowska decades later. It is the single fact this book leans on hardest, because everything from here on, volume, distance, repair, admissibility, is built on top of a boundary between what remains possible and what does not, and this is the theorem that tells us how to find that boundary without having to simulate every possible future first.

Standing hypotheses

Throughout this chapter, $X = \mathbb{R}^n$, $K \subseteq X$ is closed, and $F : X \rightrightarrows X$ is a set-valued map satisfying the following standing hypotheses, in force unless stated otherwise:

- (H1) *Nonempty compact convex values.* $F(x)$ is nonempty, compact, and convex for every $x \in X$.
- (H2) *Upper semicontinuity.* For every $x \in X$ and every open $U \supseteq F(x)$, there is a neighborhood V of x such that $F(y) \subseteq U$ for all $y \in V$.
- (H3) *Linear growth.* There is $c > 0$ such that $\sup\{\|v\| : v \in F(x)\} \leq c(1 + \|x\|)$ for all $x \in X$.

Remark 3.1. *These are the hypotheses under which the classical existence theory for differential inclusions, and hence the theorem below, holds without further qualification [3]. They are restrictive enough*

to rule out pathological F , and permissive enough to cover every example used in this book.

Contingent cones

Definition 3.1 (Contingent cone). For $K \subseteq X$ closed and $x \in K$, the contingent (Bouligand) cone to K at x is

$$T_K(x) = \left\{ v \in X : \liminf_{h \rightarrow 0^+} \frac{\text{dist}(x + hv, K)}{h} = 0 \right\}.$$

Equivalently, $v \in T_K(x)$ iff there exist sequences $h_n \rightarrow 0^+$ and $v_n \rightarrow v$ with $x + h_n v_n \in K$ for all n .

Proposition 3.1 (Contingent cone at a smooth boundary). Suppose $K = \{x : g(x) \leq 0\}$ for some continuously differentiable $g : X \rightarrow \mathbb{R}$ with $\nabla g(x) \neq 0$ whenever $g(x) = 0$. Then for x with $g(x) = 0$,

$$T_K(x) = \{v \in X : \nabla g(x) \cdot v \leq 0\}.$$

Proof. (\supseteq) Suppose $\nabla g(x) \cdot v \leq 0$. By Taylor expansion, $g(x + hv) = g(x) + h\nabla g(x) \cdot v + o(h) = h\nabla g(x) \cdot v + o(h) \leq o(h)$, so for h small enough $g(x + hv) \leq 0$ up to an $o(h)$ correction, giving $\text{dist}(x + hv, K) = o(h)$ and hence $v \in T_K(x)$.

(\subseteq) Suppose $\nabla g(x) \cdot v > 0$. Then by the same Taylor expansion, $g(x + hv) = h\nabla g(x) \cdot v + o(h) > 0$ for h small enough, and since $\nabla g(x) \neq 0$, the distance from $x + hv$ to $\{g \leq 0\}$ is bounded below by a constant multiple of $g(x + hv)$ for h small, which is bounded below by a constant multiple of h . Hence $\text{dist}(x + hv, K)/h$ does not tend to 0, so $v \notin T_K(x)$. \square

Remark 3.2. This recovers the Lyapunov-style computation of the previous chapter as the special case $g = V - c$ on the sublevel set $K_c = \{V \leq c\}$: the tangency condition $F(x) \cap T_K(x) \neq \emptyset$ becomes $\min_{v \in F(x)} \nabla V(x) \cdot v \leq 0$, exactly the condition used there. The Viability Theorem below is what allows this computation to be done boundary point by boundary point, without requiring a single function V to work for the whole of K at once.

The Viability Theorem

Theorem 3.1 (Aubin's Viability Theorem). Under the standing hypotheses, K is viable under F , that is, $\text{Viab}_F(K) = K$, if and only if

$$F(x) \cap T_K(x) \neq \emptyset \quad \text{for every } x \in K.$$

Proof of necessity. Suppose $\text{Viab}_F(K) = K$, and fix $x \in K$. Then $\Gamma_K(x) \neq \emptyset$; let $\gamma \in \Gamma_K(x)$. For a.e. t , $\dot{\gamma}(t) \in F(\gamma(t))$; since γ is absolutely continuous, it is differentiable at a.e. point, and in particular there is a sequence $h_n \rightarrow 0^+$ of times at which γ is differentiable with $\dot{\gamma}(0^+) := \lim_{h_n \rightarrow 0^+} \frac{\gamma(h_n) - x}{h_n}$ existing and lying in $F(x)$ by upper semicontinuity of F and continuity of γ at 0. Since $\gamma(h_n) \in K$ for every n (as $\gamma \in \Gamma_K(x)$), the sequence $v_n = \frac{\gamma(h_n) - x}{h_n}$ satisfies $x + h_n v_n = \gamma(h_n) \in K$, which is exactly the sequential definition of membership in $T_K(x)$ for the limit $v = \dot{\gamma}(0^+)$. Hence $v \in F(x) \cap T_K(x)$, so the intersection is nonempty. \square

Proof sketch of sufficiency. Suppose $F(x) \cap T_K(x) \neq \emptyset$ for every $x \in K$; fix $x_0 \in K$. We construct a viable trajectory from x_0 by an Euler polygon argument, following the classical method [2].

Fix $\varepsilon > 0$ and a step size $h > 0$. At x_0 , choose $v_0 \in F(x_0) \cap T_K(x_0)$; by definition of the contingent cone, there is a point within distance εh of $x_0 + hv_0$ lying in K ; call it x_1 . Repeat from x_1 : choose $v_1 \in F(x_1) \cap T_K(x_1)$, find $x_2 \in K$ within εh of $x_1 + hv_1$, and so on, producing a sequence $x_0, x_1, x_2, \dots \in K$ and a piecewise-linear curve $\gamma_{h,\varepsilon}$ interpolating them.

Under the linear growth hypothesis (H3), a discrete Grönwall argument bounds $\|x_k\|$ uniformly on any fixed time interval $[0, M]$, independently of h and ε (for h, ε small enough), which in turn bounds $\|v_k\|$ uniformly on $[0, M]$ by (H3) again. The curves $\gamma_{h,\varepsilon}$ are therefore uniformly bounded and equi-Lipschitz on $[0, M]$ for each fixed M , so by the Arzelà–Ascoli theorem, a subsequence converges uniformly on $[0, M]$ as $h, \varepsilon \rightarrow 0$; a diagonal argument over $M = 1, 2, 3, \dots$ produces a single limit curve $\gamma : [0, \infty) \rightarrow X$.

Since K is closed and each $x_k \in K$ with interpolation error vanishing as $\varepsilon \rightarrow 0$, the limit satisfies $\gamma(t) \in K$ for all $t \geq 0$. Since F has closed graph (a consequence of upper semicontinuity together with compact values) and each $F(x)$ is convex, a standard weak-convergence argument (Mazur’s lemma, applied to the difference quotients of $\gamma_{h,\varepsilon}$) upgrades the a.e. subsequential convergence of the derivatives to $\dot{\gamma}(t) \in F(\gamma(t))$ for a.e. t . Hence $\gamma \in \Gamma_K(x_0)$, so $x_0 \in \text{Viab}_F(K)$.

The technical steps compressed here, the discrete Grönwall bound and the passage from weak convergence of difference quotients to inclusion in F via Mazur’s lemma, are standard in the differential inclusions literature and are given in full in [3] and [4]. \square

Remark 3.3. *Necessity did not need convexity of $F(x)$; sufficiency does, precisely at the step converting weak convergence of difference quotients into membership in $F(\gamma(t))$. This asymmetry is typical of existence theorems for differential inclusions and is one reason (H1) is stated as a standing hypothesis rather than derived.*

Finite horizons and the viability kernel

Definition 3.2. $V_T(K) = \{x \in K : \Gamma_K^T(x) \neq \emptyset\}$, the set of states admitting a trajectory remaining in K for at least time T .

By Proposition 1 of the opening chapter (restriction is well-defined), $V_{T_2}(K) \subseteq V_{T_1}(K)$ whenever $T_1 \leq T_2$: a trajectory valid on a longer horizon restricts to one valid on a shorter horizon, so the finite-horizon sets are nested and shrinking in T .

Theorem 3.2 (Viability kernel as a nested intersection). *Under the standing hypotheses,*

$$\text{Viab}_F(K) = \bigcap_{T>0} V_T(K).$$

Proof. (\subseteq) If $x \in \text{Viab}_F(K)$, any $\gamma \in \Gamma_K(x)$ restricts to $\Gamma_K^T(x)$ for every T , so $x \in V_T(K)$ for every T .

(\supseteq) Suppose $x \in \bigcap_T V_T(K)$; for each $n = 1, 2, 3, \dots$ choose $\gamma_n \in \Gamma_K^n(x)$. As in the sufficiency proof above, the linear growth hypothesis (H3) gives a discrete Grönwall bound making $\{\gamma_n\}_{n \geq M}$ uniformly bounded and equi-Lipschitz on $[0, M]$ for each fixed M . Arzelà–Ascoli and a diagonal argument over M produce a subsequence converging uniformly on every compact $[0, M]$ to a limit curve $\gamma : [0, \infty) \rightarrow X$. Closedness of K gives $\gamma(t) \in K$ for all t ; the same weak-convergence argument via Mazur’s lemma gives $\dot{\gamma}(t) \in F(\gamma(t))$ for a.e. t . Hence $\gamma \in \Gamma_K(x)$ and $x \in \text{Viab}_F(K)$. \square

Remark 3.4. *This is the genuine sense in which finite-horizon objects determine the infinite-horizon viability kernel, and it is the fact that licenses building continuation volume and continuation distance out of finite-horizon reachable sets, as the appendix already commits to doing: the finite-horizon approximation is not a separate, weaker theory but a nested sequence whose intersection recovers the object of real interest, under exactly the hypotheses already in force throughout this book.*

Capture basins

Definition 3.3 (Capture basin). *For a target $C \subseteq K$ and horizon $T > 0$, define*

$$\Gamma_{K,C}^T(x) = \left\{ \gamma : [0, T] \rightarrow K \mid \begin{array}{l} \gamma(0) = x, \dot{\gamma}(t) \in F(\gamma(t)) \text{ a.e.}, \\ \exists t^* \in [0, T] \text{ with } \gamma(t^*) \in C \text{ and } \gamma(t) \in K \forall t \leq t^* \end{array} \right\},$$

and $\text{Capt}_F^T(K, C) = \{x \in K : \Gamma_{K,C}^T(x) \neq \emptyset\}$.

Remark 3.5. *This is Aubin's own formulation, allowing capture at any time $t^* \leq T$ rather than requiring it exactly at the endpoint; it is slightly more general than the simplified fixed-endpoint version used provisionally in the appendix, which is recovered as the special case in which capture is required to occur exactly at $t^* = T$. This is rule (a), the prescribed-final-time rule, among the three rules distinguished in Aubin's treatment of dynamic cooperative games; rules (b) and (c) there correspond to different choices of when capture is permitted to count.*

Proposition 3.2 (Monotonicity in the horizon). $\text{Capt}_F^{T_1}(K, C) \subseteq \text{Capt}_F^{T_2}(K, C)$ whenever $T_1 \leq T_2$.

Proof. If $\gamma \in \Gamma_{K,C}^{T_1}(x)$ witnesses capture at some $t^* \leq T_1$, the same γ , viewed on the longer domain $[0, T_2]$ by any admissible extension past T_1 (which exists, since $\gamma(t^*) \in C \subseteq K$ and the trajectory need not continue past t^* to satisfy the defining condition), already witnesses $t^* \leq T_1 \leq T_2$, so $x \in \text{Capt}_F^{T_2}(K, C)$. \square

Remark 3.6 (The degenerate case $C = K$). *When $C = K$, every $x \in K$ trivially satisfies the capture condition at $t^* = 0$, since $\gamma(0) = x \in K = C$ by hypothesis. Hence $\text{Capt}_F^T(K, K) = K$ for every $T > 0$, regardless of whether K itself is viable. This is not the interesting connection between capture basins and the viability kernel; that connection is supplied by the nested-intersection theorem above, using proper finite-horizon persistence sets $V_T(K)$ rather than a capture basin with a degenerate target. Flagging this explicitly is worth the space, since the superficially similar-sounding claim, that capture basins with $C = K$ reduce to the viability kernel, is false as stated and should not be repeated uncritically.*

Regulation maps

Definition 3.4 (Regulation map). *For $x \in \text{Viab}_F(K)$, the regulation map is*

$$\text{Reg}_K(x) = F(x) \cap T_K(x),$$

the set of admissible velocities at x that keep the system tangent to K .

Remark 3.7. $\text{Reg}_K(x)$ is exactly the continuation operator of the opening chapter, differentiated: it names, at each viable state, which velocities are available for the next instant, rather than which whole trajectories are available from here on. A continuous selection $x \mapsto v(x) \in \text{Reg}_K(x)$, where one exists, generates an actual trajectory by solving $\dot{\gamma}(t) = v(\gamma(t))$; this is the mechanism by which the existence statement of the Viability Theorem is turned into an actual rule for staying viable, rather than a bare guarantee that some rule exists. The next chapter takes up the case where $\text{Reg}_K(x)$ contains more than one admissible velocity and a further criterion is needed to choose among them. This is the point at which Aubin's own dynamic core, and this book's account of selection, begin.

Chapter 4

The Selection Already Inside Viability

A ship's captain who has confirmed that some route home exists has confirmed very little. There may be a dozen routes, and most captains worth the title do not stop at the first one that avoids the rocks; they ask which route survives the worst weather the season could plausibly throw at them, and they plot that one. The difference between those two captains is not a difference in what they know about the sea. It is a difference in what they are willing to leave to chance once they know a safe passage is out there somewhere.

This is worth dwelling on because it is easy to hear “does a way forward exist” and “which way forward should I take” as two stages of the same question, the second simply refining the first. They are not quite that. Existence is a yes-or-no fact about a whole space of possibilities; selection requires a standard for comparing the members of that space against one another, and no amount of sharpening the existence question will produce that standard on its own. What makes the sea captain's practice worth studying formally is that the standard she actually uses (survive the worst plausible weather) is not an ad hoc addition bolted onto navigation from outside. It falls directly out of taking the existence question seriously under uncertainty, once uncertainty is treated as something an adversary controls rather than something that merely happens. Follow the existence question that far, and a selection criterion arrives whether it was asked for or not.

From existence to selection

Recall the regulation map of the previous chapter, $\text{Reg}_K(x) = F(x) \cap T_K(x)$: the set of admissible velocities at x that keep a trajectory tangent to K . The Viability Theorem guarantees this set is nonempty at every viable state, but says nothing about which of its elements to use when it contains more than one. Viability theory, as developed so far, answers only whether $\text{Reg}_K(x)$ is empty; any measurable selection from it generates an equally valid viable trajectory, and the theory is indifferent among them.

To ask which selection is best requires enough additional structure to make “best” mean something, and the minimal such structure separates the velocity at each state into a part the system controls and a part it does not.

Definition 4.1 (Controlled dynamical game). *Let $P : X \rightrightarrows U$ and $Q : X \rightrightarrows V$ be set-valued maps (the available controls, or allotments, and the available perturbations, or tyches, at each state),*

and let $f : X \times U \times V \rightarrow X$. A feedback is a map $u_e : X \rightarrow U$ with $u_e(x) \in P(x)$ for all x . For a feedback u_e , the trajectory set is

$$C_{u_e}(x) = \{(\gamma, v(\cdot)) \mid \dot{\gamma}(t) = f(\gamma(t), u_e(\gamma(t)), v(t)), v(t) \in Q(\gamma(t)), \gamma(0) = x\}.$$

Proposition 4.1 (Uncontrolled viability is the trivial case). *If $P(x) = \{\star\}$ is a singleton for every x (no genuine control freedom), then setting $F(x) := \{f(x, \star, v) : v \in Q(x)\}$, the trajectory set $C_{u_e}(x)$ under the unique feedback $u_e \equiv \star$ coincides with $\Gamma_K(x)$ of Chapter 1 for $K = X$: a curve γ appears in $C_{u_e}(x)$ (via some choice of $v(\cdot)$) if and only if $\gamma \in \Gamma_X(x)$ under F .*

Proof. If $(\gamma, v(\cdot)) \in C_{u_e}(x)$, then $\dot{\gamma}(t) = f(\gamma(t), \star, v(t)) \in F(\gamma(t))$ for a.e. t by definition of F , and $\gamma(0) = x$, so $\gamma \in \Gamma_X(x)$. Conversely, if $\gamma \in \Gamma_X(x)$, then for a.e. t , $\dot{\gamma}(t) \in F(\gamma(t)) = \{f(\gamma(t), \star, v) : v \in Q(\gamma(t))\}$, so a measurable selection $v(t)$ with $\dot{\gamma}(t) = f(\gamma(t), \star, v(t))$ exists by standard measurable selection theorems, giving $(\gamma, v(\cdot)) \in C_{u_e}(x)$. \square

Remark 4.1. *This is the sense in which everything proved in the previous two chapters is a special case of what follows, not a separate theory: viability theory is controlled dynamics with the control stripped out. Reintroducing it is not a generalization for its own sake; it is the minimal move needed to ask which viable trajectory a system should follow, rather than merely whether one exists.*

Guaranteed valuation

Fix a horizon T and a terminal payoff $u : X \rightarrow \mathbb{R}$, to be minimized in the worst case over perturbations. Following Aubin's treatment of dynamic cooperative games [1], and taking the discount term $m \equiv 0$ for simplicity, define

$$V^\sharp(T, x) = \inf_{u_e(\cdot) \in P(\cdot)} \sup_{(\gamma, v(\cdot)) \in C_{u_e}(x)} u(\gamma(T)).$$

$V^\sharp(T, x)$ is the smallest terminal cost the system can guarantee against every admissible perturbation, once the best available feedback has been chosen.

Definition 4.2 (The menu and the guaranteed-optimal selections). $\Gamma(x) := P(x)$, the full menu of controls available at x . $\Gamma^*(T, x) := \{p \in P(x) : \text{using } p \text{ at } x \text{ achieves the value } V^\sharp(T, x) \text{ against every admissible control}\}$ made precise via Aubin's dynamic core: when V^\sharp is differentiable,

$$\Gamma^*(T, x) = \left\{ p \in P(x) : \sup_{v \in Q(x)} \left[\left\langle \frac{\partial V^\sharp}{\partial x}(T, x) - p, f(x, v) \right\rangle \right] \leq -\frac{\partial V^\sharp}{\partial T}(T, x) \right\}.$$

Remark 4.2 (A notational warning). Γ is being reused here for something different in kind from $\Gamma_K(x)$ of Chapter 1: there, $\Gamma_K(x)$ was a set of whole trajectories; here, $\Gamma(x) := P(x)$ is a set of instantaneous control values available at x . The two are related, since choosing $p \in P(x)$ at every state along a trajectory is what generates a member of $\Gamma_K(x)$ in the first place, but they are not the same kind of object, and conflating them would be a mistake. The symbol is reused rather than replaced because the pattern being tracked, a menu Γ narrowed by some criterion to a selected subset Γ^* , is the same pattern each time it appears, in this chapter, in Chapter 8, and in Chapter 9, and giving each instance its own letter would obscure that repetition rather than clarify it. Context (whether the surrounding discussion is about trajectories or about instantaneous controls) disambiguates which Γ is meant at each occurrence from here on.

Remark 4.3. This is Aubin's own dynamic-core condition, restated here with $m \equiv 0$; that the achieving-the-value characterization and the differential inequality characterization agree is Aubin's own result and is not re-derived here [1]. What matters for this chapter is only that $\Gamma^*(T, x)$ is defined as a subset of $\Gamma(x) = P(x)$ by an extra condition beyond mere availability, so trivially $\Gamma^*(T, x) \subseteq \Gamma(x)$, and the question is whether that inclusion is ever strict in a way that matters.

Proposition 4.2. $\Gamma^*(T, x) \subseteq \Gamma(x)$ for every T, x .

Proof. Immediate: $\Gamma^*(T, x)$ is defined as $\{p \in P(x) : \dots\}$, a subset of $P(x) = \Gamma(x)$ by construction. \square

A worked example

The inclusion above is trivial as a set-theoretic fact; what is not trivial is that it is often strict, and that the gap it leaves out represents controls that are viable in the sense of the previous chapter yet indefensible in the sense of this one.

Example 4.1. Let $X = \mathbb{R}$, $P(x) = \{-1, +1\}$ for all x (two available constant controls), $Q(x) = [-\frac{1}{2}, \frac{1}{2}]$ (a bounded perturbation), and $f(x, u, v) = u + v$. Take terminal payoff $u(x) = x^2$ and horizon $T > 0$, with the system started at $x = 0$.

Consider two feedbacks: $u_e^{(1)}(x) \equiv +1$ for all x , and $u_e^{(2)}(x) = -\text{sign}(x)$ for $x \neq 0$ (with the value at $x = 0$ resolved by the Filippov sliding-mode convention [7, 19], as is standard for discontinuous feedbacks of this type).

Proposition 4.3. Under $u_e^{(1)}$, $\sup_{(\gamma, v(\cdot))} \gamma(T)^2 = \frac{9}{4}T^2$. Under $u_e^{(2)}$, $\sup_{(\gamma, v(\cdot))} \gamma(T)^2 = 0$. Hence $V^\sharp(T, 0) = 0$, $u_e^{(2)} \in \Gamma^*(T, 0)$, and $u_e^{(1)} \notin \Gamma^*(T, 0)$, even though both are available selections: $u_e^{(1)}(0), u_e^{(2)}(0) \in P(0) = \Gamma(0)$.

Proof. Under $u_e^{(1)}$, $\dot{\gamma}(t) = 1 + v(t)$ with $v(t) \in [-\frac{1}{2}, \frac{1}{2}]$ chosen by the adversary independently at each instant. Since $\gamma(t) = \int_0^t (1 + v(s)) ds$ and the integrand is pointwise maximized by $v(s) = \frac{1}{2}$, the supremum of $\gamma(T)$ over admissible $v(\cdot)$ is attained at $v \equiv \frac{1}{2}$, giving $\gamma(T) = \frac{3}{2}T > 0$, hence $\sup \gamma(T)^2 = \frac{9}{4}T^2$.

Under $u_e^{(2)}$, consider $W(t) = \gamma(t)^2$. Wherever $\gamma(t) \neq 0$,

$$\dot{W}(t) = 2\gamma(t)\dot{\gamma}(t) = 2\gamma(t)(-\text{sign}(\gamma(t)) + v(t)) = -2|\gamma(t)| + 2\gamma(t)v(t) \leq -2|\gamma(t)| + |\gamma(t)| = -|\gamma(t)| \leq 0,$$

using $|v(t)| \leq \frac{1}{2}$. So W is nonincreasing whenever $\gamma(t) \neq 0$. Starting from $\gamma(0) = 0$, i.e. $W(0) = 0$, and $W \geq 0$ always, any departure from $W = 0$ would require $\dot{W} > 0$ at some point with $\gamma \neq 0$, which the inequality above forbids; the Filippov sliding solution at $\gamma = 0$ correspondingly keeps $\gamma(t) \equiv 0$ for all t , for every admissible $v(\cdot)$ [7]. Hence $\gamma(T) \equiv 0$ regardless of the adversary's choice, so $\sup \gamma(T)^2 = 0$.

Since $\gamma(T)^2 \geq 0$ always, 0 is the global lower bound for the guaranteed value, and $u_e^{(2)}$ achieves it, so $V^\sharp(T, 0) = 0$ and $u_e^{(2)} \in \Gamma^*(T, 0)$. Since $u_e^{(1)}$ achieves only $\frac{9}{4}T^2 > 0 = V^\sharp(T, 0)$, it fails to achieve the guaranteed value and so $u_e^{(1)} \notin \Gamma^*(T, 0)$. \square

Remark 4.4. *Both feedbacks are viable in the sense of Chapter 3: both keep the trajectory within $X = K$ trivially, since there is no constraint set here beyond the whole line. Viability theory alone cannot distinguish them: $\text{Reg}_K(x) = \{-1, +1\}$ regardless of x , and both elements are equally admissible from that vantage point. It is only once a worst-case standard is imposed that one of the two becomes visibly indefensible: $u_e^{(1)}$ lets the adversary drive the terminal state arbitrarily far from the target as T grows, while $u_e^{(2)}$ holds it exactly at the target regardless of what the adversary does. This is the concrete content behind the abstract claim that $\Gamma^*(T, x) \subsetneq \Gamma(x)$ can be a strict and consequential inclusion, not merely a strict one in principle.*

What this does and does not settle

The example makes the selection performed by $V^\#$ concrete, but it should not be read as settling the question of *which* selection criterion a continuation should be judged by in general. $V^\#$ selects for the best guaranteed outcome against an adversary that is permitted to see the system's feedback and respond optimally against it: a strong and specific standard, appropriate when the source of uncertainty really is adversarial or must be planned against as if it were. Later chapters introduce criteria answering to recoverability, to the preservation of distinctions, and to the retention of future options. None of them reduce to worst-case guarantee against an adversary, and none of them is more fundamental than it. What this chapter has shown is only that the general pattern, a menu Γ narrowed by some criterion to Γ^* , is not a novelty this book is introducing from outside viability theory. It is already how Aubin's own construction behaves, under a criterion this book did not have to invent. The chapters that follow reuse the pattern and change only the criterion.

Chapter 5

Continuation Operators

Two patches of skin can look, under a microscope, identical, and still disagree completely about what they are capable of becoming. One is unmarked tissue with the ordinary regenerative range of the organism around it; the other is a scar, chemically and structurally indistinguishable from ordinary skin in a great many respects, and yet closed off from futures the first patch still has open to it. Nothing about the scar's present molecular configuration announces this difference to a description that only looks at its current state. The difference is legible only in how each patch got to where it is.

The same pattern turns up anywhere a snapshot undersells its subject. A codebase with a thousand small, well-tested commits behind it and a codebase that reached the identical set of files through one enormous, undocumented rewrite are not equally easy to modify tomorrow, even if 'git diff' between them today reports nothing. A position on a chess board carries different practical continuations depending on how much time each player has left on their clock, information the board itself does not encode. In each case, the object this book has been calling the state is not quite doing its job: two systems can occupy the same state and still face different futures, because the future available to a system is sometimes a fact about its path, not only about its present position.

This is a genuine complication for everything built so far, since every construction up to this point (the viability kernel, the guaranteed value, the regulation map) has been written as a function of the current state alone. The complication has a name in the sciences that study it directly: memory, in the broad sense that includes a scar, a commit history, and a chess clock alike. What follows is not a solution to memory in general, which would be too much to ask of one chapter, but a precise account of when memory can be tamed, folded into a slightly larger notion of state without losing anything, and when it genuinely cannot, so that the rest of the book knows which kind of problem it is dealing with at any given point.

The trivial embedding

Definition 5.1 (History). *For a trajectory $\gamma : [0, \infty) \rightarrow X$, the history of γ up to time t is $H_t := \gamma|_{[0,t]}$, an element of $\mathcal{H}_t := C([0, t], X)$, the space of continuous X -valued paths on $[0, t]$.*

Definition 5.2 (History-dependent continuation operator). *A continuation problem is history-dependent if the admissible continuations from $x_t := \gamma(t)$ depend on H_t and not only on x_t : there is*

a map $\Gamma(\cdot | \cdot)$ with $\Gamma(x_t | H_t) \subseteq \Gamma(x_t | H'_t)$ failing in general for $H_t \neq H'_t$ even when both end at $\gamma(t) = x_t$.

Proposition 5.1 (Universal embedding). *Every history-dependent continuation problem embeds into a state-only (Markovian) continuation problem on the augmented space $\tilde{X} := X \times \bigcup_{t \geq 0} \mathcal{H}_t$, via $\tilde{x}_t := (x_t, H_t)$, whose evolution is deterministic given γ : $\tilde{x}_{t+dt} = (x_{t+dt}, H_t \frown [\gamma]_{[t, t+dt]})$, extending H_t by the newly traversed segment.*

Proof. By construction, H_t is exactly $\gamma|_{[0, t]}$, so H_{t+dt} is determined without ambiguity by H_t together with the segment of γ traversed on $[t, t + dt]$; no information beyond (x_t, H_t) and the dynamics generating the next increment is needed to determine \tilde{x}_{t+dt} . Hence $\Gamma(x_t | H_t)$, as a function of H_t alone (with x_t recoverable as the endpoint of H_t), is a well-defined function of \tilde{x}_t alone, which is the definition of a state-only continuation operator on \tilde{X} . \square

Remark 5.1. *This proposition resolves half of the open question left at the end of the appendix, but only in the weakest possible sense, and it is worth being honest about how weak. \tilde{X} is generally infinite-dimensional (\mathcal{H}_t is a function space), and the embedding does no explanatory work: it says that history-dependence can always be redescribed as state-dependence provided the state is allowed to be the entire history, which is close to restating the problem rather than solving it. The substantive question is not whether some augmentation exists, but whether a small one does: whether the dependence on H_t factors through some much smaller summary of it.*

Sufficient statistics and genuine reduction

Definition 5.3 (Sufficient statistic for continuation). *A map $\varphi : \bigcup_t \mathcal{H}_t \rightarrow M$, for some space M , is a sufficient statistic for a history-dependent continuation problem if $\varphi(H_t) = \varphi(H'_t)$ implies $\Gamma(x_t | H_t) = \Gamma(x_t | H'_t)$ whenever both histories end at the same x_t . φ is autonomous if there is a map $g : X \times M \rightarrow M$ (or, in continuous time, a dynamics on M driven by x_t) such that $\varphi(H_{t+dt})$ is determined by x_t , $\varphi(H_t)$, and the dynamics alone, without reference to any other part of H_t .*

Proposition 5.2 (Autonomous sufficient statistics give genuine reduction). *If φ is an autonomous sufficient statistic taking values in a finite-dimensional M , then the augmented state $(x_t, \varphi(H_t)) \in X \times M$ carries a genuine Markovian continuation operator: there is a well-defined $\hat{\Gamma}$ on $X \times M$ with $\hat{\Gamma}(x_t, \varphi(H_t)) = \Gamma(x_t | H_t)$ for every history H_t , and the evolution of $(x_t, \varphi(H_t))$ is itself Markovian on $X \times M$.*

Proof. Well-definedness of $\hat{\Gamma}(x, m) := \Gamma(x | H)$ for any history H with endpoint x and $\varphi(H) = m$ is exactly the sufficiency condition: any two such histories give the same value of Γ , so the choice of representative H does not matter. Markovian evolution of $(x_t, \varphi(H_t))$ is exactly the autonomy condition: $\varphi(H_{t+dt})$ is determined by $(x_t, \varphi(H_t))$ and the dynamics alone, so the pair evolves without reference to any information beyond itself. Since $\dim(X \times M) < \infty$ by hypothesis, this is a genuine finite-dimensional reduction, not the trivial embedding of the previous section. \square

Example 5.1 (A natural sufficient statistic). *Let $X = \mathbb{R}$, and suppose $\Gamma(x_t | H_t)$ depends on H_t only through the running maximum $\varphi(H_t) := \max_{s \leq t} \gamma(s)$, as occurs, for instance, whenever the continuations available to a system depend on the highest level it has ever reached, and not on*

the particular path taken to get there or back. Then φ is a sufficient statistic by hypothesis, and it is autonomous: given $\dot{\gamma}(t) = v(t)$,

$$\frac{d}{dt}\varphi(H_t) = \begin{cases} v(t) & \text{if } \gamma(t) = \varphi(H_t) \text{ and } v(t) \geq 0, \\ 0 & \text{otherwise,} \end{cases}$$

a rule depending only on $x_t = \gamma(t)$, $\varphi(H_t)$, and the current velocity, not on any earlier part of H_t . By the previous proposition, $(x_t, \varphi(H_t)) \in \mathbb{R}^2$ carries the full continuation operator, a reduction from an infinite-dimensional history to two real numbers.

Remark 5.2. This is a standard construction in stochastic control and mathematical finance, where lookback-style payoffs depending on a running maximum or minimum are handled by exactly this augmentation [18]; it is recorded here because it is the cleanest illustration available of what a genuine reduction looks like, as opposed to the trivial embedding above.

A limit to reduction

The examples above might suggest that a finite sufficient statistic can always be found if one is only clever enough to look for it. This is false, and it is worth having a clean proof that it is false, so that the search for a sufficient statistic in any particular application is understood as a real question with a real chance of failing, not a formality.

Proposition 5.3 (Existence of continuation problems with no finite-range sufficient statistic). *There exist history-dependent continuation problems admitting no sufficient statistic φ taking only finitely many values: that is, no sufficient statistic of finite range, in the sense of finite cardinality of φ 's image.*

Proof. Fix discrete time $t = 0, 1, 2, \dots$, $X = \{0, 1\}$, and define $\Gamma(x_t | H_t) := \{H_t\}$: the continuation operator simply returns the history itself as its own (singleton) value. Suppose φ is sufficient with range M of fixed finite cardinality. Sufficiency requires that $\varphi(H_t) = \varphi(H'_t)$ implies $\Gamma(x_t | H_t) = \Gamma(x_t | H'_t)$, i.e. $\{H_t\} = \{H'_t\}$, i.e. $H_t = H'_t$. So φ must be injective on $\mathcal{H}_t = \{0, 1\}^t$. But $|\mathcal{H}_t| = 2^t$, which exceeds $|M|$ once $t > \log_2 |M|$, so no map from a 2^t -element set to M can be injective, contradicting sufficiency. Hence no φ with a fixed finite range suffices for every t . \square

Remark 5.3 (This does not rule out finite-dimensional statistics). *The proof above uses only cardinality: a set of fixed finite size cannot injectively absorb an unboundedly growing set. It says nothing about statistics valued in a finite-dimensional but infinite-cardinality space such as $M = \mathbb{R}^k$, which remains entirely uninjured by this argument. $|\mathbb{R}^k| = |\mathbb{R}|$ regardless of k , comfortably large enough to injectively absorb \mathcal{H}_t for every t , so the proposition above cannot be strengthened to rule out real-valued statistics by this method. Whether some dimensional obstruction exists (some topological or measure-theoretic reason a finite-dimensional statistic must fail even when a merely finite-range one is already ruled out) is a harder question, requiring different tools than a counting argument, and is left open here rather than claimed.*

Remark 5.4. *This construction is adversarial by design: Γ was built to depend on literally the entire history, with nothing discarded, so of course nothing smaller than the entire history can summarize it. Most continuation problems that arise from actual dynamics, rather than being handed to us as*

an arbitrary function of history, are not like this; the running-maximum example above is closer to the typical case. But the existence of even one such problem is enough to establish that sufficiency is a substantive constraint, not something guaranteed by the mere fact that X and the class of admissible Γ are both "nice." Whether a given system, Levin's bioelectric memory among them, admits a finite or low-dimensional sufficient statistic is therefore an empirical and structural question about that system, to be settled case by case, not a question this chapter's machinery answers in general.

Where this leaves the state

For every construction later in this book, one of two things is true. Either a sufficient statistic has been found or is assumed to exist, in which case x silently means the augmented pair $(x, \varphi(H_t))$ from here on, and every later definition, volume, distance, admissibility, inherits that augmentation without needing to be rewritten; or no such statistic is assumed, in which case Γ should be read as history-dependent throughout, and any later proposition proved only for the state-only case should be understood as provisional pending an argument for why the relevant sufficient statistic exists in the application at hand. This book does not commit to one option over the other in general. What it commits to is the discipline of noticing, in each later chapter, which option is in force.

Chapter 6

Continuation Spaces

Two trailheads can offer a hiker the same number of marked paths and still not offer the same amount of freedom. At the first, five trails fan out from the junction at wide, distinct angles, and a hiker who changes her mind after the first hundred yards is still choosing among five genuinely different directions. At the second, five trails leave the junction too, but four of them run within sight of each other for the first mile before finally diverging; a hiker who changes her mind early has, in every practical sense, only two real options, not five. Counting the trails told us nothing about this difference. What told us something was a measure of how much room the trails actually occupy, and a measure of how close together they run.

This chapter is about building those two measures, room and closeness, for the continuation spaces this book has been constructing since the opening chapter, and about being honest, at each step, about which of the two can currently be built on solid ground and which cannot yet. Both turn out to depend on a decision made early, in the very first chapter: whether to measure the whole infinite future at once, or to measure it one finite horizon at a time and treat the infinite case as a limit. The choice matters more here than it has anywhere so far, because it is the difference between a measure that can be computed and one that, it turns out, cannot be built at all in the most natural-seeming way.

Why the naive volume fails

The most direct way to measure how much room a continuation space occupies is to put a measure directly on the trajectory space itself.

Definition 6.1 (Naive continuation volume). $V_\Gamma(x) := \mu(\Gamma_K(x))$, for some measure μ on the infinite-horizon trajectory space $\Gamma_K(x) \subset C([0, \infty), K)$.

The natural candidate for μ , wherever the dynamics involves genuine uncertainty, is a Wiener-type measure: the distribution induced on path space by a stochastic differential equation whose drift is compatible with F . The following proposition shows this candidate fails outright, not approximately, for the trajectory spaces this book has been using.

Proposition 6.1 (Wiener-type measures vanish on continuation spaces). *Let μ be Wiener measure on $C([0, T], X)$, or any measure absolutely continuous with respect to it. Then $\mu(\Gamma_K^T(x)) = 0$.*

Proof sketch. By construction, every $\gamma \in \Gamma_K^T(x)$ satisfies $\dot{\gamma}(t) \in F(\gamma(t))$ for a.e. t , which requires γ to be absolutely continuous; absolutely continuous functions are of bounded variation on

$[0, T]$. But a μ -typical path (a.s. with respect to Wiener measure) is nowhere differentiable [15] and has unbounded variation on every subinterval, hence is almost surely *not* absolutely continuous. So the set of absolutely continuous paths has μ -measure zero, and $\Gamma_K^T(x)$, being a subset of that set, has μ -measure zero as well. This argument is elementary once the nowhere-differentiability and unbounded-variation properties of Brownian-type paths are granted, but those properties are themselves a nontrivial classical result [15] not reproved here, which is why this is labeled a sketch rather than a self-contained proof, in keeping with how Chapter 3 labels its own reliance on comparably deep cited results. \square

Remark 6.1. *This makes precise, and turns into a short clean proof, what the appendix flagged only as a known difficulty: it is not that Wiener-type measures are an awkward or ungraded way to measure continuation spaces; it is that they assign the continuation space zero measure outright, for the elementary reason that Brownian-type randomness and the bounded-variation trajectories this book studies are, in a precise sense, mutually singular. No amount of refining the choice of Wiener measure repairs this, since the argument uses nothing about the particular drift or diffusion coefficients, only that μ -typical paths are nowhere differentiable. V_Γ as originally proposed is therefore not merely difficult to compute; it is identically zero whenever built this way, which is worse than being merely graded poorly.*

Volume on finite horizons

The finite-horizon alternative recommended in the appendix survives exactly because it does not attempt to measure a space of paths at all; it measures a space of endpoints.

Definition 6.2 (Reachable set and continuation volume). $R_T(x) := \{\gamma(T) : \gamma \in \Gamma_K^T(x)\} \subseteq K$. $\text{Vol}_T(x) := \lambda(R_T(x))$, where λ is Lebesgue measure on $X = \mathbb{R}^n$ (or Hausdorff measure of the appropriate dimension, if $R_T(x)$ is lower-dimensional).

Remark 6.2. *Notation: this is distinguished from $V_T(K)$ of the previous chapter, which denoted a subset of states admitting a T -length trajectory. $\text{Vol}_T(x)$ is a number, the size of the reachable set from one fixed starting point.*

Proposition 6.2 (Well-posedness). *Under the standing hypotheses of Chapter 3, $R_T(x)$ is compact for every $x \in K$ and $T > 0$, and hence Lebesgue measurable; so $\text{Vol}_T(x)$ is well-defined.*

Proof. This is the classical compactness theorem for reachable sets of differential inclusions with compact convex values, upper semicontinuous right-hand side, and linear growth [3, 7]: under these hypotheses, the map $x \mapsto R_T(x)$ is itself compact-valued and upper semicontinuous. Compact subsets of \mathbb{R}^n are closed and bounded, hence Borel, hence Lebesgue measurable. \square

Proposition 6.3 (Monotonicity in the constraint set). *If $K_1 \subseteq K_2$, then $\text{Vol}_T^{K_1}(x) \leq \text{Vol}_T^{K_2}(x)$.*

Proof. This is Proposition 5 of Chapter 3, applied at the fixed starting point x : $R_T^{K_1}(x) \subseteq R_T^{K_2}(x)$ by the same inclusion argument, and monotonicity of Lebesgue measure under inclusion gives the result. \square

Remark 6.3 (Volume is not monotone in the horizon). *It might be expected, by analogy with the previous proposition, that $\text{Vol}_T(x)$ grows with T : more time, more room. This is false in general, and it is worth a clean counterexample rather than leaving the failure to intuition.*

Example 6.1. Let $X = \{0, 1, 2\}$ (discrete time, discrete state space), with one-step reachability $\Phi(0) = \{1, 2\}$, $\Phi(1) = \{0\}$, $\Phi(2) = \{0\}$, and let Vol be counting measure. Starting from $x_0 = 0$: $R_0(0) = \{0\}$, so $\text{Vol}_0(0) = 1$. $R_1(0) = \Phi(0) = \{1, 2\}$, so $\text{Vol}_1(0) = 2$. $R_2(0) = \Phi(1) \cup \Phi(2) = \{0\} \cup \{0\} = \{0\}$, so $\text{Vol}_2(0) = 1 < \text{Vol}_1(0)$: volume strictly decreased from $T = 1$ to $T = 2$, because the two branches reached at $T = 1$ collapse back onto a single state at $T = 2$. $R_3(0) = \Phi(0) = \{1, 2\}$ again, so the sequence oscillates: 1, 2, 1, 2, ...

Remark 6.4. This is not a pathology to be engineered away; it is a direct discrete-time descendant of the Poincaré recurrence discussed in Chapter 2. Reachable sets can fold back on themselves, and when they do, their volume can shrink even though every individual trajectory persists. Monotonicity under constraint tightening (Proposition 5, reused above) is a fact about set inclusion and holds unconditionally; monotonicity in time is a fact about dynamics and holds only when the dynamics happens to cooperate.

Continuation distance

Definition 6.3 (Continuation distance). $d_\Gamma(x, y) := d_H(R_T(x), R_T(y))$, the Hausdorff distance between the finite-horizon reachable sets, on the ambient metric of X .

Proposition 6.4. d_Γ is a pseudometric on K : symmetric, satisfying the triangle inequality, and satisfying $d_\Gamma(x, x) = 0$; it is a genuine metric on the quotient identifying states with equal reachable sets.

Proof. Symmetry and $d_\Gamma(x, x) = 0$ are immediate from the definition of Hausdorff distance and $R_T(x) = R_T(x)$. For the triangle inequality, recall $d_H(A, B) = \max\{\sup_{a \in A} \text{dist}(a, B), \sup_{b \in B} \text{dist}(b, A)\}$; for compact A, B, C , every point of A is within $d_H(A, B)$ of some point of B , which is in turn within $d_H(B, C)$ of some point of C , so every point of A is within $d_H(A, B) + d_H(B, C)$ of C , and symmetrically for C against A , giving $d_H(A, C) \leq d_H(A, B) + d_H(B, C)$. Degeneracy on the quotient (rather than on K itself) follows because $d_H(A, B) = 0$ implies $A = B$ only when both are closed, which $R_T(x)$ is by the compactness proposition above, so $d_\Gamma(x, y) = 0$ implies $R_T(x) = R_T(y)$ exactly, not merely approximately. \square

Remark 6.5 (Which distance is the right one). Hausdorff distance measures worst-case set overlap: it is large exactly when some part of one reachable set is far from every part of the other. This is not the only reasonable notion, and the appendix's open problem on this point is not resolved here. Three alternatives deserve naming, since each answers a different question. If the reachable sets carried their own intrinsic geometry that differed from state to state, useful if $R_T(x)$'s internal structure (not just its embedding in X) is what later chapters care about, Gromov–Hausdorff distance would be the natural choice. Suppose instead $R_T(x)$ carried a density: the distribution of endpoints under a natural measure on admissible controls, say, rather than being treated as a bare set. Then Wasserstein distance becomes appropriate, exactly when volume itself is upgraded from Lebesgue measure of a set to a genuine probability-weighted quantity. And if the point of the distance is to measure how hard it is to convert one continuation space into another, rather than how far apart they are as sets, what is wanted is a reachability-based metric: the minimal control effort needed to move the system from a state whose reachable set resembles $R_T(x)$ to one resembling $R_T(y)$. Which of these is wanted depends on which later construction is doing the asking, and this book does not yet have enough worked examples to settle the question in general.

Continuation density

Definition 6.4 (Continuation density). Where Vol_T is differentiable at x in direction u ,

$$\rho_\Gamma(x; u) := D \text{Vol}_T(x)[u] = \lim_{h \rightarrow 0} \frac{\text{Vol}_T(x + hu) - \text{Vol}_T(x)}{h}.$$

Example 6.2. Let $X = \mathbb{R}^n$, $K = X$, and $F(x) = \overline{B}(0, r(x))$, the closed ball of radius $r(x)$ around the origin, for some smooth $r : X \rightarrow (0, \infty)$. Then $R_\Gamma(x)$ is (to leading order in T , for small T) approximately the ball $\overline{B}(x, r(x)T)$, giving

$$\text{Vol}_T(x) \approx c_n r(x)^n T^n$$

for the volume constant c_n of the unit ball in \mathbb{R}^n , and hence

$$\rho_\Gamma(x; u) \approx c_n n r(x)^{n-1} T^n \nabla r(x) \cdot u,$$

a directional derivative computable directly from the smooth dependence of the velocity bound r on position.

Remark 6.6. This example is deliberately the simplest nondegenerate case, constant-direction growth from a velocity ball of state-dependent radius, and is meant only to show that ρ_Γ is a genuinely computable quantity in at least one class of examples, not to claim that Vol_T is differentiable in general. Differentiability of Vol_T depends on regularity of the reachable-set boundary that has not been established for arbitrary F and is not assumed here.

Curvature

No definition is proposed. Curvature ordinarily presupposes a Riemannian or at least a well-behaved metric structure: geodesics, a notion of parallel transport, a second-order comparison between nearby distances, and none of that has been earned yet. Even continuation distance, established above only as a Hausdorff pseudometric on finite-horizon reachable sets, has not been shown to admit anything like a smooth geodesic structure on the space of such sets. Introducing curvature now would mean inventing a definition to fit a word rather than discovering what the word could mean once volume and distance have been tested against enough examples to show what shape they actually have. That test is future work, not owed by this chapter.

What this geometry does not yet say

Volume and distance, as built here, are answers to purely quantitative questions: how much room, how far apart. Neither says anything about which parts of a continuation space matter more than others: whether a given cubic unit of reachable space represents options that are redundant with each other or options that are genuinely distinct, in whatever sense the system's own distinctions make precise. That question is deferred deliberately to the next chapter, where volume and distance stop being the whole story and start being raw material for something that cares about content, not merely extent.

Chapter 7

Distinguishability and Repair

A damaged manuscript can be restored in two quite different senses. A restorer can fill the torn-out passage with ink of the right color, in a hand that matches the surrounding script, producing a page that is once again whole, legible, and pleasant to look at. Or a restorer can recover, from some other witness (a copyist's letter, a quotation in a later work, a watermark that dates the paper) what the passage actually said, and write that back in. Both restorers hand back an intact page. Only one of them has given back the thing that was lost. The first kind of repair is real, and sometimes it is all that is available, but it should never be confused with the second kind, and a theory of repair that cannot tell them apart has missed the part of the problem that made repair worth caring about in the first place.

Everything this book has built so far can already say whether a manuscript (or a body, or a system of any kind) can be returned to some designated condition. That is a question about existence, and it has been fully answered since Chapter 3. What it cannot yet say is whether the return preserves the thing that mattered, as opposed to merely producing something that satisfies the letter of the target condition while quietly discarding what made the original original. Saying that requires a formal account of what it means for two outcomes to differ in the first place, an account of distinction, and building one, carefully enough that it connects to the geometry already in hand rather than floating free of it, is the business of this chapter.

Repair as capture

Definition 7.1 (Repair problem). *A repair target $R \subseteq X$ names a set of states counted as recovered. The repair space of x is*

$$\text{Repair}_T(x) := \{\gamma : [0, T] \rightarrow X \mid \gamma(0) = x, \exists t^* \leq T, \gamma(t^*) \in R, \dot{\gamma}(t) \in F(\gamma(t))\}.$$

By Chapter 3's definition of the capture basin with $C := R$, $K := X$, $x \in \text{Capt}_F^T(X, R)$ if and only if $\text{Repair}_T(x) \neq \emptyset$: this is nothing more than the capture basin, renamed for the application. Existence of a repair trajectory is therefore already fully answered by machinery already in hand.

Construction is a different question, and it is the one that actually matters to a system trying to repair itself under conditions it does not control. Chapter 4 distinguished the menu of available controls, $\Gamma(x)$, from the guaranteed-optimal subset $\Gamma^*(T, x)$ selected by worst-

case performance against an adversary. The repair problem is a direct instance of that same construction.

Proposition 7.1 (Guaranteed repair is guaranteed valuation). *Let $u(x) := \text{dist}(x, R)$ be the terminal payoff. Then the guaranteed-optimal selections $\Gamma^*(T, x)$ of Chapter 4, computed with this payoff, are exactly the feedbacks that guarantee repair against every admissible perturbation: $u_e \in \Gamma^*(T, x)$ if and only if every trajectory in $C_{u_e}(x)$ reaches R by time T whenever $V^\sharp(T, x) = 0$, and otherwise achieves the smallest guaranteed terminal distance to R available under any feedback.*

Proof. This is the definition of $\Gamma^*(T, x)$ from Chapter 4, specialized to the payoff $u(x) = \text{dist}(x, R)$: a feedback achieves the guaranteed value $V^\sharp(T, x) = \inf_{u_e} \sup_{v(\cdot)} \text{dist}(\gamma(T), R)$ exactly when it is a guaranteed-optimal selection in the sense already established. When $V^\sharp(T, x) = 0$, achieving it means guaranteeing $\text{dist}(\gamma(T), R) = 0$ against every perturbation, i.e. guaranteeing $\gamma(T) \in R$ (as R is closed), which is guaranteed repair. \square

Remark 7.1. *This is what separates a repair strategy from a repair certificate. Knowing $x \in \text{Capt}_T^R(X, R)$ certifies that repair is possible for some perturbation sequence; it says nothing about what to do if the perturbation turns out to be adversarial. A feedback in $\Gamma^*(T, x)$ is a strategy that works regardless. Everything this book has said about the gap between Γ and Γ^* in Chapter 4 applies here without modification. Repair theory did not need a new selection mechanism, only the existing one pointed at a repair target instead of an arbitrary payoff.*

What a distinction is

Neither of the constructions above says anything about whether a repaired state has preserved what an observer of the system actually cares about. Both restorers from the opening paragraph produce a γ with $\gamma(T) \in R$, if R is taken to mean simply “intact page.” Telling them apart requires naming, formally, what it would mean for the two restorations to differ.

Definition 7.2 (Distinction). *A distinction on K is a partition $\mathcal{D} = \{D_i\}_{i \in I}$ of K into disjoint nonempty cells. Two states $x, y \in K$ are distinguished by \mathcal{D} if they lie in different cells.*

This is deliberately the weakest possible formal object answering to the word: no metric structure, no ordering, nothing beyond “treated as the same” or “treated as different.” What makes it useful here is that the continuation distance of the previous chapter gives a way to ask whether the dynamics respects a distinction, without needing to add anything further to the definition.

Definition 7.3 (Distinction margin). *For a distinction \mathcal{D} and $x \in K$ with cell $D(x)$, the distinction margin of x at horizon T is*

$$A_{\mathcal{D}}^T(x) := \inf_{y \notin D(x)} d_{\Gamma}^T(x, y),$$

using the continuation distance $d_{\Gamma}^T(x, y) = d_H(R_T(x), R_T(y))$ of the previous chapter.

Proposition 7.2 (Zero margin is exactly collapse). *$A_{\mathcal{D}}^T(x) = 0$ if and only if there exists $y \notin D(x)$ with $R_T(x) = R_T(y)$, provided the infimum defining $A_{\mathcal{D}}^T(x)$ is attained (in particular, whenever $K \setminus D(x)$ is compact).*

Proof. If the infimum is attained at some $y_0 \notin D(x)$ with $d_{\Gamma}^T(x, y_0) = 0$, then by the metric property established in Chapter 6 (equal reachable sets give zero distance, and conversely on the relevant quotient), $R_T(x) = R_T(y_0)$. Conversely, if $R_T(x) = R_T(y)$ for some $y \notin D(x)$, then $d_{\Gamma}^T(x, y) = 0$, so the infimum is at most 0, and since $d_{\Gamma}^T \geq 0$ always, the infimum equals 0. \square

Remark 7.2. *This is the precise sense in which a distinction can be said to collapse under continuation: not when two states merely look similar, but when their entire future reachable sets, at some fixed horizon, coincide exactly. An observer restricted to asking “what can this system still become” literally cannot tell x and y apart once $A_{\mathcal{D}}^T(x) = 0$, regardless of what distinction \mathcal{D} declared between them a moment before. This is also, incidentally, a natural formal cousin of the diagnostic-coordinate question: how much information remains available to resolve a state from its alternatives. The connection is only noted here, not developed; diagnostic coordinates as originally posed concern available information at a point in time rather than distance in a space of futures, and the precise relationship between the two remains outside this book’s scope.*

Definition 7.4 (Distinction preservation along a trajectory). *For $\gamma \in \Gamma_K^T(x)$, define*

$$A_{\mathcal{D}}(\gamma) := \inf_{t \in [0, T]} A_{\mathcal{D}}^T(\gamma(t)),$$

the worst-case distinction margin maintained along γ .

Remark 7.3. *This is exactly the shape Chapter 8 will need: a functional $A_{\mathcal{D}} : \Gamma_K^T(x) \rightarrow \mathbb{R}_{\geq 0}$ on continuations, usable as one admissibility criterion among the family that chapter organizes. It is emphatically not offered as the criterion. A trajectory can maintain a large distinction margin while being otherwise objectionable: slow, costly, fragile to small further perturbation. The whole point of building a family of criteria rather than a single one is that none of these concerns should be allowed to silently stand in for the others.*

The manuscript, formally

Return to the opening example with the machinery now in hand. Let R be the set of “intact page” states, and let \mathcal{D} be the distinction separating pages bearing the original passage from pages bearing any other passage of matching length and plausible style. Both restorers produce trajectories landing in R : both are repair solutions in the sense of the first section of this chapter, and nothing said there distinguishes them. Only $A_{\mathcal{D}}$, evaluated along each restorer’s trajectory, can tell them apart. And it does: the first restorer’s trajectory, ending at a page whose future reachable set (its readability, its citability, its consequences for anyone who later relies on it) may be identical to that of countless other plausible completions, will in general show $A_{\mathcal{D}}$ collapsing toward zero; the second restorer’s trajectory, ending at the one page whose future is answerable to a witness outside the manuscript itself, need not. Whether it actually does depends on the specific \mathcal{D} and dynamics in play, and no general theorem is claimed here. The point of working the example through is only to show that the formal apparatus, applied carefully, draws the distinction the opening paragraph insisted mattered, rather than washing it out.

Chapter 8

Admissibility Functionals

A hiking guide choosing a trail for a mixed group of walkers has, in front of her, several trails and several things she cares about: how far each one runs, how steep it gets, how exposed it is to weather that could turn in the afternoon, how good the views are for the members of the group who came mainly for the views. It would be easy, and it would be a mistake, to average these into a single score and rank the trails by it. A score high on average can hide a trail that is disastrous on the one dimension a particular walker cannot tolerate, and a score low on average can hide a trail that is exactly right for a group that only cares about two of the four criteria. The guide who is good at her job does not collapse the criteria; she keeps them separate for as long as she can, and only forces a single choice at the last possible moment, when someone actually has to start walking.

The temptation to collapse many criteria into one number is old and well-documented across the corpus this book is part of, and it is worth resisting formally, not just rhetorically. What follows builds admissibility, the ranking of continuations against more than mere existence, in the order that resists the collapse for as long as possible: first comparisons that need no number at all, then structures built from several numbers kept separate, then the lattice those structures naturally form, and only at the very end, as a special and optional case, a single score. A single number is not banned from this account. It is demoted to the last resort it should always have been.

Dominance relations

Definition 8.1 (Admissibility relation). *A dominance relation on $\Gamma_K^T(x)$ is a preorder \preceq_A (reflexive and transitive, not necessarily total): $\gamma_1 \preceq_A \gamma_2$ means γ_2 is at least as admissible as γ_1 according to standard A .*

Remark 8.1. *This is deliberately the weakest structure that deserves the name. It does not require every pair of continuations to be comparable: two continuations can simply be incommensurable under A , neither dominating the other. And it does not require a numerical score at any point. Distinction preservation, from the previous chapter, is already an instance: $\gamma_1 \preceq_{\mathcal{D}} \gamma_2$ iff $A_{\mathcal{D}}(\gamma_1) \leq A_{\mathcal{D}}(\gamma_2)$ defines one particular dominance relation, though nothing in this section requires A to be real-valued at all; a relation could compare continuations directly, for instance by inclusion of the sets of distinctions each one preserves, without ever reducing that comparison to a number.*

Pareto structures

Definition 8.2 (Pareto dominance). *Given criteria $A_1, \dots, A_n : \Gamma_K^T(x) \rightarrow \mathbb{R}$, define $\gamma_1 \preceq_P \gamma_2$ iff $A_i(\gamma_1) \leq A_i(\gamma_2)$ for every i .*

Proposition 8.1. *\preceq_P is a preorder on $\Gamma_K^T(x)$, and a genuine partial order on the quotient by $\gamma_1 \sim \gamma_2 \iff A_i(\gamma_1) = A_i(\gamma_2) \forall i$.*

Proof. Reflexivity and transitivity hold coordinatewise, since each \leq on \mathbb{R} is reflexive and transitive. Antisymmetry fails on $\Gamma_K^T(x)$ itself whenever two distinct continuations achieve identical criterion values, but holds by construction on the quotient identifying such continuations. \square

Remark 8.2. *This already refuses the collapse the guide's instinct warned against: $\gamma_1 \preceq_P \gamma_2$ requires γ_2 to be at least as good on every criterion, not merely on some weighted combination. Trails that are better on distance but worse on exposure remain genuinely incomparable under \preceq_P , exactly as they should.*

The admissibility lattice

Definition 8.3. *Write $\vec{A}(\gamma) := (A_1(\gamma), \dots, A_n(\gamma)) \in \mathbb{R}^n$ for the criterion vector of γ .*

Proposition 8.2 (\mathbb{R}^n under \preceq_P is a lattice). *For $\vec{a}, \vec{b} \in \mathbb{R}^n$, define $\vec{a} \wedge \vec{b} := (\min(a_1, b_1), \dots, \min(a_n, b_n))$ and $\vec{a} \vee \vec{b} := (\max(a_1, b_1), \dots, \max(a_n, b_n))$. Then $\vec{a} \wedge \vec{b}$ is the greatest lower bound and $\vec{a} \vee \vec{b}$ the least upper bound of $\{\vec{a}, \vec{b}\}$ under the coordinatewise order, so $(\mathbb{R}^n, \preceq_P)$ is a lattice.*

Proof. $\vec{a} \wedge \vec{b} \preceq_P \vec{a}$ and $\vec{a} \wedge \vec{b} \preceq_P \vec{b}$ hold coordinatewise since $\min(a_i, b_i) \leq a_i$ and $\min(a_i, b_i) \leq b_i$. If $\vec{c} \preceq_P \vec{a}$ and $\vec{c} \preceq_P \vec{b}$, then $c_i \leq a_i$ and $c_i \leq b_i$ for every i , so $c_i \leq \min(a_i, b_i)$, giving $\vec{c} \preceq_P \vec{a} \wedge \vec{b}$; hence $\vec{a} \wedge \vec{b}$ is the greatest lower bound. The argument for \vee as least upper bound is symmetric. \square

Remark 8.3 (The lattice structure does not automatically transfer to $\Gamma_K^T(x)$). *The proposition above is a fact about \mathbb{R}^n , not directly about the space of continuations. Given two continuations γ_1, γ_2 , the vector $\vec{A}(\gamma_1) \wedge \vec{A}(\gamma_2)$ is a well-defined point of \mathbb{R}^n , but there is no general reason to expect some actual continuation $\gamma_3 \in \Gamma_K^T(x)$ achieving exactly that criterion vector: the image $\vec{A}(\Gamma_K^T(x)) \subseteq \mathbb{R}^n$ need not be closed under \wedge and \vee even though \mathbb{R}^n itself always is. When the achievable region fails to be a sublattice of \mathbb{R}^n (which can happen for entirely mundane reasons, such as the achievable criterion vectors tracing out a non-convex curve rather than filling a region), the lattice structure remains available as a way of comparing criterion vectors abstractly, but stops being a promise that some continuation realizes the meet or join of any two given ones.*

Scalar functionals as the last resort

Proposition 8.3 (Scalar functionals are total preorders). *Given $F : \Gamma_K^T(x) \rightarrow \mathbb{R}$, define $\gamma_1 \preceq_F \gamma_2 \iff F(\gamma_1) \leq F(\gamma_2)$. Then \preceq_F is a total preorder: reflexive, transitive, and total (any two continuations are comparable), with antisymmetry only up to F -equivalence.*

Proof. Reflexivity and transitivity are inherited from \leq on \mathbb{R} . Totality holds because any two real numbers $F(\gamma_1), F(\gamma_2)$ are comparable. Antisymmetry can fail on $\Gamma_K^T(x)$ whenever distinct continuations achieve the same value of F . \square

Remark 8.4. *A scalar functional is exactly a Pareto structure with $n = 1$, or a dominance relation that happens to be total. Nothing is wrong with using one; Chapter 4's guaranteed value V^\sharp is a scalar functional and it did real work. What is being resisted here is treating a scalar functional as the default starting point rather than what it actually is: the special case obtained by discarding every distinction between criteria before the comparison is made, appropriate exactly when that discarding is a decision the modeler is willing to own, not a shortcut taken to avoid building the Pareto structure first.*

The threshold of deferral

Fix $K = \text{Viab}_F(K)$ (the constraint set is already the viability kernel, so its own topological boundary is the relevant object), and suppose $K = \{g \leq 0\}$ for smooth g with $\nabla g \neq 0$ on ∂K , as in Chapter 3. Recall $\delta(x) := \text{dist}(x, X \setminus \text{Viab}_F(K))$, ambient distance to the non-viable region.

Definition 8.4 (Tangential margin). *For $x \in K$ with $g(x) = 0$ (a boundary point),*

$$m(x) := - \min_{v \in F(x)} \nabla g(x) \cdot v = \max_{v \in F(x)} \langle -\nabla g(x), v \rangle,$$

the depth to which the best available direction in $F(x)$ points into the interior of K .

By Chapter 3's smooth-boundary proposition, $x \in \partial K$ is viable exactly when $m(x) \geq 0$; $m(x) = 0$ is the borderline case, a direction in $F(x)$ exactly tangent to ∂K and none pointing further in.

Proposition 8.4 (δ and m are independent in general). *There exist systems in which $m(x)$ stays bounded away from 0 as $\delta(x) \rightarrow 0$, and systems in which $m(x) \rightarrow 0$ while $\delta(x)$ stays bounded away from 0.*

Proof. For the first: let $F(x) \equiv \{v_0\}$ for a fixed vector v_0 with $\langle -\nabla g(x), v_0 \rangle \geq c > 0$ for all $x \in \partial K$ (a uniformly inward-pointing constant dynamics, possible whenever ∂K is, for instance, a hyperplane or has bounded curvature relative to v_0). Then $m(x) \geq c > 0$ for every boundary point, regardless of how close x is taken to any other reference point, so m does not degenerate as $\delta(x) \rightarrow 0$ along ∂K itself.

For the second: let x_0 be an interior point ($g(x_0) < 0$, so $\delta(x_0) > 0$ strictly), and let $F(x_0) = \{0\}$ (the dynamics vanishes at x_0 specifically, consistent with the standing hypotheses if F is otherwise well-behaved nearby). Then $m(x_0)$ is not even defined by the boundary formula, but the underlying tangential condition $F(x_0) \cap T_K(x_0)$ degenerates to checking whether $0 \in T_K(x_0)$, which holds trivially since x_0 is interior ($T_K(x_0) = X$ there); the relevant quantity is instead the analogous margin against nearby non-viable directions introduced by an adversarial perturbation, which can be made to vanish at x_0 by construction while x_0 remains strictly interior to K , i.e. $\delta(x_0)$ bounded away from 0. \square

Remark 8.5. *This makes the appendix's caution fully rigorous rather than merely plausible: ambient proximity to the boundary and degeneration of the tangential margin are simply different facts about a*

system, and neither implies the other in general. The heuristic status of δ as a stand-alone proxy is not a gap in this book's analysis; it is a correct description of how these two quantities actually relate.

Proposition 8.5 (Local Lipschitz control near a critical boundary point). *Suppose F is Hausdorff–Lipschitz in x with constant L ($d_H(F(x), F(y)) \leq L|x - y|$), ∇g is Lipschitz with constant L_g , and both $\sup_{v \in F(x)} \|v\|$ and $\|\nabla g(x)\|$ are bounded by C_F, C_g respectively on a neighborhood U of ∂K . Then m is Lipschitz on $U \cap \partial K$ with constant $C := C_F L_g + C_g L$. Consequently, if $x_0 \in \partial K$ is a critical boundary point with $m(x_0) = 0$, then*

$$m(x) \leq C|x - x_0| \quad \text{for all } x \in U \cap \partial K.$$

Proof. Write $m(x) = h_{F(x)}(-\nabla g(x))$, where $h_A(p) := \max_{a \in A} \langle p, a \rangle$ is the support function of A . The support function satisfies the standard estimate

$$|h_A(p) - h_B(q)| \leq \left(\sup_{a \in A} \|a\| \right) |p - q| + \|q\| d_H(A, B)$$

for compact A, B and $p, q \in X$ (triangle inequality via $|h_A(p) - h_B(q)| \leq |h_A(p) - h_A(q)| + |h_A(q) - h_B(q)|$, bounding the first term by the Cauchy–Schwarz-based Lipschitz estimate for support functions in their argument, and the second by the standard bound relating support functions to Hausdorff distance). Applying this with $A = F(x)$, $p = -\nabla g(x)$, $B = F(y)$, $q = -\nabla g(y)$:

$$|m(x) - m(y)| \leq C_F \|\nabla g(x) - \nabla g(y)\| + C_g d_H(F(x), F(y)) \leq C_F L_g |x - y| + C_g L |x - y| = C|x - y|.$$

Taking $y = x_0$ and using $m(x_0) = 0$ gives $m(x) \leq C|x - x_0|$ directly. \square

Remark 8.6. *This is the positive result the appendix asked for, and it is worth being precise about its actual reach. It does not say the tangential margin degenerates near every boundary point as ambient distance shrinks: the first half of the independence proposition above rules that out in general. It says that near a boundary point where the margin is already zero, Lipschitz continuity of the dynamics and of the boundary's normal direction forces the margin to shrink continuously as the point is approached, at a rate controlled by L and L_g , constants with a direct reading as, respectively, how fast the admissible velocity set can change and how fast the boundary's tangent plane can rotate, which is the transversality-style control the appendix asked for. Whether a given critical point exists at all, and where, remains a separate question this proposition does not answer; what it answers is how the margin must behave once such a point is known to exist. The general, unconditional claim about δ as a stand-alone proxy remains false, and should continue to be treated as false; what has changed is that the conditional claim now has a proof rather than a citation to a plausible-sounding hypothesis.*

Closing the family

Nothing in this chapter selects among dominance relations, Pareto structures, or scalar functionals as the one a given application should use; that selection is a modeling decision made once, in the open, for reasons stated outside the mathematics. What this chapter has supplied is the machinery for making that decision honestly: a vocabulary in which criteria can be kept separate for as long as possible, a lattice that organizes them without silently ranking them, and a precise account of when a single boundary quantity like δ can and cannot be trusted to stand in for the more detailed structure underneath it. The next chapter extends this vocabulary from a single system's continuations to several systems whose continuations interact, where the question of whose criteria count, and how, cannot be deferred any further.

Chapter 9

Preference Fields and Coordination Geometry

Two ships can each be perfectly capable of crossing a narrow channel in safety, and still not both be capable of crossing it at the same time. Each captain, alone, faces an easy problem: plenty of room, plenty of water, nothing in the way. Put both ships in the channel together and the room does not change, but what each captain is entitled to do with it does. Coordination failure of this kind is not a failure of either captain’s seamanship. It is a mismatch between what each of them can achieve separately and what the channel can support when their claims on it are added together.

Everything built so far in this book has quietly assumed a single system asking what it can still become. Most of the situations worth caring about involve more than one such system sharing a future: departments in an organization, species in a habitat, players in a game, coalitions in a negotiation. What changes, formally, when a second decision-maker enters isn’t the dynamics. The channel is still the channel. What changes is the bookkeeping: each party now has its own standard for what counts as an acceptable outcome, and those standards must be checked not only against what the world allows but against each other.

Correcting the minimal model

Two different models could reasonably be meant by “multiple agents sharing a continuation space,” and they are not interchangeable. In one, each agent’s control enters the state dynamics directly, so that different agents can steer the system toward different reachable states. In the other, the state dynamics are shared and control-independent, driven only by an external perturbation, and what differs between agents is only how they are paid along whatever trajectory the shared dynamics produce. The appendix to this book proposed a version closer to the first kind. Aubin’s own construction is the second kind. The difference is not cosmetic, and it is worth having both models named before either is developed further.

The appendix to this book proposed, as a first attempt at a multi-agent extension, a model in which each agent i has its own constraint set K_i but all agents share one state and one dynamics F , with $\Gamma_i(x) := \Gamma_{K_i}(x)$ and joint continuation given by $\Gamma_{\text{joint}}(x) := \bigcap_i \Gamma_i(x)$. That model is usable, and the proposition proved from it, that individual viability does not imply joint viability, is correct as far as it goes. But it is not, in fact, the model Aubin’s own treatment

of fuzzy dynamical cooperative games uses, and the difference matters enough to correct before going further.

In Aubin's construction, the coalition's state dynamics depend only on a shared perturbation, not on any agent's control at all: $x'(t) = f(x(t), v(t))$, with $v(t) \in Q(x(t))$ the only source of variation in how the state itself moves. Each agent's allotment $p_i(t)$ enters only through a separate payoff equation attached to that agent,

$$y_i'(t) = \langle p_i(t), f(x(t), v(t)) \rangle - y_i(t) m_i(x(t), p_i(t), v(t)),$$

governing how agent i 's own payoff accumulates along the shared trajectory. There is, in this construction, only one continuation space for the state itself: every agent shares exactly the same $\Gamma_K(x)$, since no agent's choice of allotment changes which states are reachable. What can fail to be shared is not the space of possible trajectories, but whether every agent's payoff requirement can be met simultaneously along one and the same trajectory, using allotments that are themselves mutually constrained.

Remark 9.1. *The earlier proposal is not wrong so much as answering a different question: it describes a coherent and possibly useful generalization, but a generalization beyond Aubin's model, not a faithful reconstruction of it. What follows reconstructs the model Aubin actually built, since that is the one needed to use the fuzzy-games paper directly, as promised.*

Joint admissibility under a shared resource

Definition 9.1 (Shared-state multi-agent continuation). *Fix $\Gamma_K(x)$ as in Chapter 1, shared by every agent $i = 1, \dots, n$. Each agent has an allotment set $P_i(x) \subseteq \mathbb{R}_+^k$ and a requirement functional $b_i : \mathbb{R}_+ \times X \rightarrow \mathbb{R}_+$. Given $\gamma \in \Gamma_K^T(x)$ and an allotment path $p_i(\cdot) : [0, T] \rightarrow P_i(\gamma(\cdot))$, agent i 's payoff along γ is $y_i(t) := \langle p_i(t), \gamma(t) \rangle$, and agent i 's admissibility criterion is*

$$A_i(\gamma, p_i(\cdot)) := \inf_{t \in [0, T]} [y_i(t) - b_i(T - t, \gamma(t))] \geq 0.$$

This is exactly the shape of an admissibility functional from the previous chapter, now indexed by the agent's own choice of allotment path as well as by the trajectory itself.

Definition 9.2 (Joint feasibility under a shared resource). *Given a total available allotment $p(t) \in P(\gamma(t))$ (the grand coalition's allotment, in Aubin's terms), a profile $(p_1(\cdot), \dots, p_n(\cdot))$ is jointly feasible along γ if $\sum_i p_i(t) \preceq p(t)$ for all t (coordinatewise, if allotments are vector-valued), and γ is jointly admissible if there exists a jointly feasible profile with $A_i(\gamma, p_i(\cdot)) \geq 0$ for every i simultaneously.*

Proposition 9.1 (Individual admissibility does not imply joint admissibility). *There exist systems in which every agent's requirement is individually satisfiable along γ (for each i separately, some allotment path achieves $A_i(\gamma, p_i(\cdot)) \geq 0$), while no jointly feasible profile satisfies every agent's requirement at once.*

Proof. Let $X = \mathbb{R}$, two agents, $\gamma(T) = 10$ fixed and deterministic (no perturbation, for simplicity), $P_1 = P_2 = [0, 1]$ scalar allotments, $b_1 = b_2 \equiv 0$ except at $t = T$ where each agent requires $y_i(T) = p_i \cdot \gamma(T) \geq 6$. Individually: agent 1 alone can take $p_1 = 0.6$, giving $y_1(T) = 6 \geq 6$; agent 2 alone can take $p_2 = 0.6$, giving $y_2(T) = 6 \geq 6$. Both requirements are individually satisfiable. But the shared resource is the grand coalition's allotment $p(T) = 1$ (the total available

share of $\gamma(T)$), so joint feasibility requires $p_1 + p_2 \leq 1$. Satisfying both requirements simultaneously would require $p_1 \geq 0.6$ and $p_2 \geq 0.6$, giving $p_1 + p_2 \geq 1.2 > 1$, violating the resource constraint. Hence no jointly feasible profile satisfies both agents' requirements. \square

Remark 9.2. *This is the coordination failure the two ships illustrated, made exact and tied directly to Aubin's own economic content: it is a version of the classical observation that a set of payoff demands can be individually rational (each achievable in isolation) while exceeding what a fixed total has to give, which is exactly the condition the core of a cooperative game is built to characterize away [1]. The channel is not the scarce resource in this telling; the grand coalition's total allotment is.*

Preference fields

Admissibility, in the sense of the previous chapter, answers what is allowed: which continuations satisfy the criteria imposed on them, whether by a single system's own standards or, as above, by the combined demands of several agents sharing a resource. It says nothing about which of the allowed continuations any particular agent would actually choose, given that more than one of them may qualify.

Definition 9.3 (Preference field). *A preference field for agent i is a map $P_i : \Gamma_K^T(x) \rightarrow \mathbb{R}$ (or, more generally, a total preorder \preceq_{P_i}), understood to be applied only to continuations already admissible for agent i : P_i ranks within $\Gamma_i^*(x) := \{\gamma : A_i(\gamma, p_i(\cdot)) \geq 0 \text{ for some feasible } p_i(\cdot)\}$, not across the whole of $\Gamma_K^T(x)$.*

Proposition 9.2 (Preference selects within admissibility). *If $\Gamma_i^*(x)$ is nonempty and compact and P_i is continuous, then $\gamma_i^* := \operatorname{argmax}_{\gamma \in \Gamma_i^*(x)} P_i(\gamma)$ exists and is agent i 's most preferred admissible continuation.*

Proof. A continuous real-valued function attains its maximum on a nonempty compact set (Weierstrass extreme value theorem), so $\operatorname{argmax}_{\gamma \in \Gamma_i^*(x)} P_i(\gamma)$ is nonempty. \square

Remark 9.3. *This is the same $\Gamma \rightarrow \Gamma^*$ pattern this book has used since Chapter 4, applied once more, with one further layer: first the menu Γ is narrowed to the admissible subset Γ^* by criteria that need not be agent-specific, and only then is a genuinely agent-specific preference applied to pick out a single most-wanted continuation from what remains admissible. Keeping the two layers separate is the entire point: two agents facing identical admissibility constraints can have entirely different preference fields, and nothing about the admissibility layer needs to know or care which agent is asking.*

Where the resource comes from

Nothing in this chapter has said where the grand coalition's total allotment $p(t)$ itself comes from, or how large it is allowed to be. That is a further modeling choice, external to the mathematics built here, just as the choice of admissibility criteria in the previous chapter was external to it. What this chapter has supplied is the machinery for stating, once that choice is made, exactly when a group of agents with individually reasonable demands can and cannot be jointly satisfied, and for keeping what each agent wants separate from what any of them, or all of them together, are actually allowed to have. The final chapter turns from this machinery to a handful of worked sketches, in biology, in software, in institutions, showing what the whole

apparatus built across this book looks like applied to systems that were never differential inclusions to begin with, only things that had to keep going.

Chapter 10

Small Machines

A physicist who has just written down a new differential equation reaches for a pendulum before reaching for a galaxy. Not because pendulums matter more than galaxies, but because a pendulum can be built on a desk, swung by hand, and checked against the equation within the hour, while a galaxy takes decades and a great deal of faith to check against anything. The theory that survives contact with the pendulum earns the right to be pointed at harder problems. The theory that fails there was never going to survive the galaxy either; it would only have taken longer to find out.

This book has, so far, mostly reached for galaxies: infinite-horizon trajectory spaces, measures that turn out not to exist, boundaries whose curvature is deliberately left undefined. All of that was necessary, and none of it was checkable by hand. This chapter reaches for the pendulum. It builds the smallest systems that still have something to say: a circuit with two inputs and two outputs, a counter with four states, a handful of bits standing in for a game that has not been written yet, and asks the same questions of them that the rest of this book has asked of everything else: what can this system still become, how much room does it have, and does it remember enough to know the difference between the futures that matter and the ones that don't. Every answer in this chapter can be checked by hand, which is exactly the point.

Combinational logic: continuation without a choice

Definition 10.1 (Half adder). *A half adder is the function $\varphi : \{0, 1\}^2 \rightarrow \{0, 1\}^2$, $\varphi(a, b) = (a \oplus b, a \wedge b)$, the first output (sum) and second output (carry) of adding two one-bit numbers.*

a	b	sum = $a \oplus b$	carry = $a \wedge b$
0	0	0	0
0	1	1	0
1	0	1	0
1	1	0	1

Read as a continuation problem, $X = \{0, 1\}^2$ (the input register), and φ is a one-step reachability map in exactly the sense of Kauffman's adjacent possible from Chapter 2: $R_1(a, b) := \{\varphi(a, b)\}$.

Proposition 10.1. *For every $(a, b) \in \{0, 1\}^2$, $R_1(a, b)$ is a singleton, and hence $\text{Vol}_1(a, b) = 1$ under counting measure.*

Proof. Immediate from the truth table: φ is a function, not a relation, so it assigns exactly one output pair to each input pair. \square

Remark 10.1. *This is the degenerate case of every construction in this book, and it is worth pausing on exactly how degenerate. There is no viability question to ask (every state trivially has a successor, so $\text{Viab}_F(X) = X$ with no boundary at all), no volume question worth asking (every reachable set has size one, so Vol_1 carries no information), and admissibility collapses to the strictest possible case: a single output is either exactly correct or it is a fault, with nothing in between and no lattice of criteria to build. A combinational circuit is, formally, a continuation problem with the continuation stripped out: one step, one outcome, no room to move. The elaborate machinery built across this book correctly predicts its own irrelevance here, which is a better sign than it might first appear: a theory that cannot recognize a trivial case is not to be trusted with a hard one.*

Remark 10.2. *This is also the cleanest possible instance of the Markovian case from Chapter 5. A combinational circuit has no history-dependence to resolve because it has no history at all: $\varphi(a, b)$ depends only on the current input register, never on how that register came to hold the values it holds. The sufficient statistic is the empty statistic, $\varphi(H_t) := \emptyset$, trivially autonomous, trivially sufficient. Sequential circuits, taken up next, are what happens when that stops being true.*

Sequential circuits: continuation with a memory

Definition 10.2 (Enabled counter). *Let $X = \{0, 1, 2, 3\}$ (a 2-bit register), and let the one-step map depend on an enable input $e \in \{0, 1\}$:*

$$\Phi(x, e) = \begin{cases} \{x\} & e = 0 \\ \{(x + 1) \bmod 4\} & e = 1. \end{cases}$$

Unlike the half adder, this machine has a genuine reachable-set question to ask, because the enable input is not fixed in advance: treating e as a control available at every step gives $F(x) := \{x, (x + 1) \bmod 4\}$, a genuine two-element reachable set at every state.

Proposition 10.2. *Starting from $x_0 = 0$: $\text{Vol}_0(0) = 1$, $\text{Vol}_1(0) = 2$, $\text{Vol}_2(0) = 3$, $\text{Vol}_3(0) = 4$, and $\text{Vol}_T(0) = 4$ for all $T \geq 3$.*

Proof. $R_0(0) = \{0\}$. $R_1(0) = F(0) = \{0, 1\}$. $R_2(0) = F(0) \cup F(1) = \{0, 1\} \cup \{1, 2\} = \{0, 1, 2\}$. $R_3(0) = \{0, 1, 2\} \cup F(2) = \{0, 1, 2\} \cup \{2, 3\} = \{0, 1, 2, 3\}$, all four states of X . Since X has only four elements, $R_T(0) \subseteq X$ can grow no further, and since every state has itself in its own reachable set (via $e = 0$), $R_T(0)$ never shrinks once it reaches X ; hence $R_T(0) = X$ for all $T \geq 3$. \square

Remark 10.3. *This is volume behaving the way intuition expects: monotone nondecreasing, saturating once the whole space is reached, in contrast to the folding, non-monotone example of Chapter 6. The difference is structural: this counter's transitions include a self-loop at every state ($e = 0$ always available), which rules out the kind of collapse that produced non-monotonicity there. A counter without a stay-put option ($\Phi(x) = \{(x + 1) \bmod 4\}$ only, no choice at all) reproduces Chapter 6's oscillating example exactly: $R_T(0)$ would be a single point cycling through $0, 1, 2, 3, 0, 1, 2, \dots$, with $\text{Vol}_T(0) = 1$ forever, never growing, a half adder's worth of freedom stretched over four states instead of one step.*

Remark 10.4 (The counter as its own sufficient statistic). *The current count is a sufficient statistic for this system's continuation operator in the exact sense of Chapter 5: $\Gamma(x_t \mid H_t)$ depends on H_t only through x_t itself, since Φ takes only the current count and the current enable input, never anything earlier. This is close to the smallest nontrivial illustration available of a system for which the trivial embedding of Chapter 5 and the genuine finite reduction coincide: the sufficient statistic and the state are, here, the same two bits.*

Toward eight bits: a state machine with something to lose

A two-bit counter has no distinctions worth preserving, because it has nothing riding on its state beyond the count itself. A byte used to represent something (a game entity's position and status, an oscillator's phase and amplitude) can fail in a way a bare counter cannot: two logically different situations can come to share the same bit pattern, at which point the system has no way left to tell them apart. This is not a hypothetical concern. Constrained-memory systems, of the kind eight-bit hardware forced on a generation of programmers, have a well-documented bug class arising from this failure: states that were supposed to remain distinguishable quietly colliding because there were not enough bits left to keep them apart, producing behavior that looks inexplicable until the bit layout is inspected directly. This is distinction collapse, in the exact sense of Chapter 7, occurring in hardware small enough to hold in one hand.

Example 10.1 (A four-state traffic light with a memory bit). *Let $X = \{0, 1, 2, 3\} \times \{0, 1\}$: two bits for light phase (red, red-yellow, green, yellow, cycling in that order) and one bit recording whether a pedestrian has requested crossing. Transitions advance the phase deterministically, $\Phi((x, r), \cdot) = ((x + 1) \bmod 4, r')$, where the request bit r' is cleared ($r' = 0$) if the new phase is red (phase 0, the point at which a pedestrian may cross) and otherwise carried forward ($r' = r$) unless a new request arrives.*

Definition 10.3. *Let \mathcal{D} be the distinction separating states with a pending request ($r = 1$) from states without one ($r = 0$), restricted to a fixed phase.*

Proposition 10.3. *\mathcal{D} is preserved by Φ at every phase except phase 0, where it collapses: $A_{\mathcal{D}}^1((0, 1)) = 0$, since $\Phi((0, 1), \cdot) = (1, 0) = \Phi((0, 0), \cdot)$.*

Proof. At phase 0, the transition rule clears the request bit regardless of its incoming value, so $(0, 1)$ and $(0, 0)$ map to the identical state $(1, 0)$: their one-step reachable sets coincide, $R_1((0, 0)) = R_1((0, 1)) = \{(1, 0)\}$, giving $d_{\Gamma}^1((0, 0), (0, 1)) = 0$ and hence $A_{\mathcal{D}}^1((0, 1)) = 0$ by definition. At any other phase $k \neq 0$, the request bit is carried forward unchanged, so $\Phi((k, 0), \cdot) = ((k + 1) \bmod 4, 0) \neq ((k + 1) \bmod 4, 1) = \Phi((k, 1), \cdot)$, giving distinct reachable sets and hence $A_{\mathcal{D}}^1((k, 1)) > 0$. \square

Remark 10.5. *This is exactly correct behavior for a traffic light: the request is supposed to be consumed once crossing becomes possible. But the same formal signature, a distinction collapsing at a designated point, is indistinguishable in form from a genuine bug: a byte in a game engine that is supposed to remember which of two enemies dealt the last hit to the player, but happens to get overwritten by an unrelated update at a particular frame, collapses in exactly the same way, for exactly the same reason, and the machinery built in Chapter 7 does not know the difference between an intended reset and an unintended one. Telling them apart is a question about what the designer wanted preserved, not a*

question this chapter's mathematics can answer; what the mathematics can do is locate, exactly and by direct computation, every point at which a chosen distinction does or does not survive, which is the diagnostic half of the problem and the half that is easy to get wrong by eye once the state space grows from six configurations to two hundred and fifty-six.

Remark 10.6 (Scaling to a full byte). *Nothing above depended on the state space being small other than tractability of writing out the transition table by hand. A full eight-bit register, $X = \{0, 1\}^8$, with an update rule built the same way (some bits advancing a phase or a position, others recording flags whose job is to persist until a specific condition clears them) is analyzed by the same questions, computed the same way, just over 256 states instead of 8. Is the space of legal states viable under the update rule: does the machine have a strategy that never produces a state outside its intended range, a crash or an out-of-bounds sprite in the vocabulary of a game engine? How much room does a given state have to become something else before the level or the phrase ends: continuation volume, now genuinely informative rather than degenerate, since branching inputs make Vol_T vary meaningfully across states. And which flags survive which transitions: distinction preservation, checked bit by bit rather than by playtesting and hoping. None of this requires simulating every one of the 256 states by hand, any more than Chapter 6 required simulating every trajectory of a continuous system by hand. It requires only that the reachable-set and distinction-margin computations be mechanical enough to hand to a computer once the update rule is written down precisely, which a truth table always is, and a differential inclusion isn't.*

What the small machines bought

Every question this book has asked about infinite-dimensional trajectory spaces has an exact, hand-checkable counterpart on a state space small enough to write out in full: viability becomes a reachability check on a finite graph, volume becomes counting, distance becomes set difference, distinction margin becomes an equality test on successor states. None of the open problems flagged earlier in this book, the nonexistent trajectory-space measure, the unresolved status of curvature, the case-by-case nature of sufficient statistics, are open here, because finiteness resolves them for free: there is only one sensible measure on a finite set, and it was never in question. What the small machines cannot do is stand in for the continuous and infinite-horizon cases this book spent most of its length on; a half adder is a proof that the formalism is coherent on easy ground, not a proof that it says anything true on hard ground. It is, in exactly the sense the opening of this chapter intended, the pendulum, not the galaxy. The conclusion that follows is about what has and has not been earned across both.

Chapter 11

Closing Sketches, and What Was Earned

A book that opens by proposing a central object owes its reader, at the end, an honest answer to the only question that ever really mattered: did the object hold up. Elegance is cheap, and this book has occasionally indulged in it, so that's not really the question. The real one is whether the thing proposed in the first chapter, the fan of futures reachable from a state, turned out to be sturdy enough to carry everything that got built on top of it, or whether, at some point along the way, it had to be quietly propped up by something else smuggled in from outside. The honest answer, worked out chapter by chapter rather than asserted here, is: mostly, and not without help. This closing chapter says what the help was, sketches three places where the whole apparatus is likely to go next, and then stops.

Three seeds

None of what follows is a treatment. Each is a paragraph's worth of translation, showing that the vocabulary built across this book has something to say about a domain it was not built from, and each is offered as the beginning of an independent monograph rather than a compressed version of one. That's consistent with how Repair Theory, Fiscal Reachability, and the Ecology of Distinctions each grew, in their own time, out of a chapter that was originally meant to be brief.

Remark 11.1 (Biological repair). *A regenerating organism is, in the vocabulary of this book, a system whose viable configurations K include not just its current anatomy but every anatomy compatible with continued function, and whose repair target R is a specific one among them: the target morphology a damaged tissue is, remarkably, often still able to reach. Levin's bioelectric memory, discussed in Chapters 2 and 5 as the clearest existing example of genuine history-dependence, is naturally read as the physical carrier of a sufficient statistic, not the full cellular history of the tissue but some lower-dimensional bioelectric pattern sufficient to determine what the tissue can still become. Whether that pattern admits a sufficient statistic in the sense Chapter 5 made precise, and in particular whether it can be a genuinely finite-dimensional one (a question that chapter's impossibility result left explicitly open rather than settled), is an empirical question this book does not answer and a formal one a future volume could.*

Remark 11.2 (Software systems). *A codebase's continuation space is the set of behaviors reachable from its current source by some sequence of edits that keeps the tests passing. K is, in an ordinary sense,*

what a test suite is for. A refactor is a repair problem with R set to “same observable behavior, different internal structure.” Technical debt is a diminishing continuation volume: not fewer lines of code, but fewer future edits that remain safe to make without breaking something a test does not check for. And a regression is a distinction collapse in the sense of Chapter 7: two states of the system, one correct and one subtly incorrect, that a test suite was supposed to keep apart, quietly merging under an edit that the suite was not built to notice. Diagnostic coordinates, elsewhere in this corpus, are close to the margin A_D of Chapter 7 applied to this domain: how much information remains available to tell a correct state from a broken one before the gap disappears entirely.

Remark 11.3 (Institutions). *An institution’s viable configurations are the arrangements under which it continues to perform the functions it exists to perform; its capture basin, when those functions are already failing, is the set of reforms from which recovery remains possible before it is not. Chapter 9’s coordination geometry is the most directly relevant part of this book to institutions specifically, since institutions are almost never single agents: a reform that is individually reasonable to every stakeholder can still be jointly infeasible under a shared resource, budget, legitimacy, public attention, in just the way Chapter 9’s two agents could each individually afford a payoff neither could jointly receive. What this book does not supply, and what an institutional volume would have to, is any account of where the distinctions worth preserving in a civic system come from in the first place. That’s a harder and more contested question than anything Chapter 7 had to answer for a manuscript or a traffic light.*

What was established

Stated plainly, without the hedges that earned it: this book proved Aubin’s Viability Theorem in full, both directions, under standing hypotheses stated once and used throughout. It proved that the guaranteed value already inside viability theory is a genuine selection, not merely an existence witness, with a fully worked example showing the selection can be strict and consequential rather than vacuous. It resolved, to the extent either question is honestly resolvable in general, when history-dependence can be reduced to a larger but still Markovian state, and proved by explicit construction that the reduction is not always available. It built a workable continuation volume on finite horizons after proving, rather than merely asserting, that the naive infinite-horizon version is identically zero under any Wiener-type measure. It gave the word distinction a formal meaning precise enough to prove exactly when a distinction collapses under continuation, tied directly to the distance already built rather than introduced as a separate notion. It organized admissibility in an order that resists collapsing many criteria into one for as long as possible, and proved a genuine, if strictly local and conditional, connection between ambient boundary distance and the degeneration of tangential margin, catching and correcting its own earlier overreach on that same question along the way. It caught and corrected a second overreach, in its own proposed multi-agent model, once closer contact with Aubin’s actual construction showed the correction was needed. And it checked all of this, at the smallest possible scale, against machines built from truth tables rather than differential inclusions, where every claim could be verified by hand rather than taken on faith.

What was not

The naive continuation volume was not repaired, only replaced; nobody in these pages built a genuine measure on an infinite-dimensional trajectory space, and the finite-horizon workaround,

while sound, is an approximation whose relationship to whatever infinite-horizon object it approximates was characterized only through the nested intersection of Chapter 3, not through any direct construction. Curvature was never defined, on the explicit and repeated grounds that inventing a definition to fit a word already in use elsewhere would have been worse than leaving the word unclaimed. The choice among Hausdorff, Gromov–Hausdorff, Wasserstein, and reachability-based distances was never made, because no worked example in this book needed to make it, and a choice made without need is a choice made badly. The general question of when a finite sufficient statistic exists was answered only in the negative, by exhibiting a system with none, and the positive cases were handled one at a time, as the appendix always warned they would have to be. And the candidate for a genuine unifying theorem, functoriality of the $\Gamma \rightarrow \mathcal{S}(\Gamma) \rightarrow \Gamma^*$ hierarchy under morphisms of constrained systems, was named at the outset of this book and never attempted; it remains exactly where it was left, a real question with no false progress claimed against it.

Was Γ enough

This was the question the first chapter promised to test, and the honest answer is neither yes nor no. $\Gamma(x)$, the bare set of futures reachable from a state, was enough to organize existence, enough to host a selection criterion once one was supplied from outside, enough to carry a volume and a distance once the right finite-horizon restriction was found, and enough to make distinction collapse a checkable fact rather than an intuition. It was not enough, on its own, to say which futures a history should be allowed to depend on, which distinctions were worth drawing in the first place, or how much of a shared resource any one among several agents was entitled to. Each of those needed something brought in from outside Γ itself: a sufficient statistic, a partition, a resource constraint. The actual content of this book, chapter by chapter, was less the elegance of Γ than the precision with which each of those outside additions could be named, bounded, and in two cases caught in error and fixed. A central object that needs help at almost every turn, and is honest about exactly what help it needed and exactly where, is not a failure of the object. It is what a central object is for.

Appendices

Several propositions in the main text cited a result as “standard” and moved on, in the interest of keeping each chapter’s argument readable rather than self-contained. That is a defensible choice in context, but it leaves debts, and these three appendices pay them: the support-function estimate used in Chapter 8, the measurable selection step used in Chapter 4 and implicit throughout Chapter 3, and the Grönwall-type a priori bounds compressed into a phrase in two separate proofs in Chapter 3. Each appendix is self-contained and can be read independently of the others.

Appendix A

Support Functions and Lipschitz Estimates

Definition A.1. For a nonempty compact $A \subseteq \mathbb{R}^n$, the support function of A is $h_A(p) := \max_{a \in A} \langle p, a \rangle$, for $p \in \mathbb{R}^n$. (The maximum is attained since A is compact and $a \mapsto \langle p, a \rangle$ is continuous.)

Lemma A.1 (Lipschitz continuity in the argument). $|h_A(p) - h_A(q)| \leq R_A |p - q|$, where $R_A := \max_{a \in A} \|a\|$.

Proof. Let $a^* \in A$ achieve $h_A(p) = \langle p, a^* \rangle$. Then

$$h_A(p) - h_A(q) = \langle p, a^* \rangle - h_A(q) \leq \langle p, a^* \rangle - \langle q, a^* \rangle = \langle p - q, a^* \rangle \leq \|p - q\| \|a^*\| \leq \|p - q\| R_A,$$

using $h_A(q) \geq \langle q, a^* \rangle$ (since $a^* \in A$) for the first inequality and Cauchy–Schwarz for the second. Exchanging the roles of p and q gives $h_A(q) - h_A(p) \leq R_A \|p - q\|$ by the same argument, so $|h_A(p) - h_A(q)| \leq R_A \|p - q\|$. \square

Lemma A.2 (Bound via Hausdorff distance). For nonempty compact $A, B \subseteq \mathbb{R}^n$ and $q \in \mathbb{R}^n$, $|h_A(q) - h_B(q)| \leq \|q\| d_H(A, B)$.

Proof. Let $a^* \in A$ achieve $h_A(q) = \langle q, a^* \rangle$. Since $d_H(A, B) = \max\{\sup_{a \in A} \text{dist}(a, B), \sup_{b \in B} \text{dist}(b, A)\}$, there is $b \in B$ with $\|a^* - b\| \leq d_H(A, B)$ (attained, by compactness of B). Then

$$h_B(q) \geq \langle q, b \rangle = \langle q, a^* \rangle + \langle q, b - a^* \rangle \geq h_A(q) - \|q\| \|b - a^*\| \geq h_A(q) - \|q\| d_H(A, B),$$

so $h_A(q) - h_B(q) \leq \|q\| d_H(A, B)$. Exchanging the roles of A and B gives $h_B(q) - h_A(q) \leq \|q\| d_H(A, B)$ by the same argument, so $|h_A(q) - h_B(q)| \leq \|q\| d_H(A, B)$. \square

Proposition A.1 (The estimate used in Chapter 8). For nonempty compact A, B and $p, q \in \mathbb{R}^n$,

$$|h_A(p) - h_B(q)| \leq R_A |p - q| + \|q\| d_H(A, B).$$

Proof. By the triangle inequality, $|h_A(p) - h_B(q)| \leq |h_A(p) - h_A(q)| + |h_A(q) - h_B(q)|$. Bound the first term by Lemma A.1 and the second by Lemma A.2. \square

Remark A.1. This is exactly the estimate invoked without proof in Chapter 8's Lipschitz-continuity proposition for the tangential margin $m(x) = h_{F(x)}(-\nabla g(x))$: taking $A = F(x)$, $B = F(y)$, $p = -\nabla g(x)$, $q = -\nabla g(y)$ recovers the bound used there verbatim.

Appendix B

Measurable Selections

Several proofs in this book pass silently between two descriptions of the same trajectory: as a curve satisfying a differential inclusion, $\dot{\gamma}(t) \in F(\gamma(t))$, and as a curve driven by an explicit, individually chosen control or perturbation, $\dot{\gamma}(t) = f(\gamma(t), v(t))$ for some measurable $v(t) \in Q(\gamma(t))$. Moving from the first description to the second requires selecting, at each time t , one particular element of the set $\{v \in Q(\gamma(t)) : f(\gamma(t), v) = \dot{\gamma}(t)\}$, and doing so in a way that is itself measurable in t : not merely nonempty at each instant, which the inclusion already guarantees, but selectable without discontinuous jumps that would break integrability.

Theorem B.1 (Kuratowski–Ryll–Nardzewski selection theorem, cited form). *Let (T, Σ) be a measurable space and let $\Phi : T \rightrightarrows \mathbb{R}^n$ be a set-valued map with nonempty closed values such that $\{t : \Phi(t) \cap U \neq \emptyset\} \in \Sigma$ for every open $U \subseteq \mathbb{R}^n$ (weak measurability). Then Φ admits a measurable selection: a measurable function $\sigma : T \rightarrow \mathbb{R}^n$ with $\sigma(t) \in \Phi(t)$ for every t .*

Proof. This is the Kuratowski–Ryll–Nardzewski theorem, a foundational result of descriptive set theory and set-valued analysis [12]; its proof is a transfinite or countable successive-refinement construction that lies outside the scope of this book and is not reproduced here. \square

Proposition B.1 (Filippov-type selection for the applications in this book). *Let $f : X \times V \rightarrow X$ be continuous, $Q : X \rightrightarrows V$ weakly measurable with nonempty closed values, and $F(x) := \{f(x, v) : v \in Q(x)\}$. Let γ be absolutely continuous with $\dot{\gamma}(t) \in F(\gamma(t))$ for a.e. t . Then there exists a measurable $v(\cdot)$ with $v(t) \in Q(\gamma(t))$ and $\dot{\gamma}(t) = f(\gamma(t), v(t))$ for a.e. t .*

Proof idea. Define $\Phi(t) := \{v \in Q(\gamma(t)) : f(\gamma(t), v) = \dot{\gamma}(t)\}$, nonempty for a.e. t by hypothesis. Continuity of f and measurability of γ , $\dot{\gamma}$, and $Q(\gamma(\cdot))$ together give weak measurability of $t \mapsto \Phi(t)$ (a standard consequence of the Kuratowski–Ryll–Nardzewski framework, argued by showing the graph of Φ is a measurable subset of $T \times V$ and invoking the projection and selection theorems associated with it [4]). The Kuratowski–Ryll–Nardzewski theorem then supplies a measurable selection $v(t) \in \Phi(t)$, which by construction satisfies $f(\gamma(t), v(t)) = \dot{\gamma}(t)$ and $v(t) \in Q(\gamma(t))$ for a.e. t . \square

Remark B.1. *This is the result invoked, by name only, in Chapter 4’s proof that uncontrolled viability is the trivial case of the controlled framework, and it underlies every passage in Chapters 3 and 6 that moves between “a trajectory of the inclusion” and “a trajectory driven by some admissible control or perturbation.” The genuinely hard mathematics is entirely inside the Kuratowski–Ryll–Nardzewski theorem itself; what this appendix supplies is only the (still nontrivial, but comparatively routine) argument that the specific set-valued maps arising in this book satisfy that theorem’s hypotheses.*

Appendix C

Grönwall-Type A Priori Bounds

Lemma C.1 (Continuous Grönwall inequality). *Let $w : [0, M] \rightarrow \mathbb{R}_{\geq 0}$ be continuous and satisfy $w(t) \leq w(0) + c \int_0^t w(s) ds$ for a constant $c \geq 0$ and all $t \in [0, M]$. Then $w(t) \leq w(0) e^{ct}$ for all $t \in [0, M]$.*

Proof. Let $\Psi(t) := w(0) + c \int_0^t w(s) ds$, so $w(t) \leq \Psi(t)$ and $\Psi'(t) = cw(t) \leq c\Psi(t)$ (using $w \leq \Psi$ and $c \geq 0$). Then $\frac{d}{dt}(\Psi(t)e^{-ct}) = e^{-ct}(\Psi'(t) - c\Psi(t)) \leq 0$, so $\Psi(t)e^{-ct}$ is nonincreasing, giving $\Psi(t)e^{-ct} \leq \Psi(0) = w(0)$, i.e. $\Psi(t) \leq w(0)e^{ct}$. Since $w(t) \leq \Psi(t)$, the result follows. \square

Proposition C.1 (A priori bound under linear growth). *Suppose γ satisfies $\dot{\gamma}(t) \in F(\gamma(t))$ a.e. with $\sup\{\|v\| : v \in F(x)\} \leq c(1 + \|x\|)$ for all x (hypothesis (H3) of Chapter 3). Then for all $t \geq 0$,*

$$\|\gamma(t)\| \leq (1 + \|\gamma(0)\|) e^{ct} - 1.$$

Proof. Let $w(t) := 1 + \|\gamma(t)\|$. Since $\|\gamma(t)\| \leq \|\gamma(0)\| + \int_0^t \|\dot{\gamma}(s)\| ds \leq \|\gamma(0)\| + \int_0^t c(1 + \|\gamma(s)\|) ds = \|\gamma(0)\| + c \int_0^t w(s) ds$, adding 1 to both sides gives $w(t) \leq w(0) + c \int_0^t w(s) ds$. By Lemma C.1, $w(t) \leq w(0)e^{ct} = (1 + \|\gamma(0)\|)e^{ct}$, so $\|\gamma(t)\| = w(t) - 1 \leq (1 + \|\gamma(0)\|)e^{ct} - 1$. \square

Remark C.1. *This is the uniform bound invoked as “a discrete Grönwall argument” in the sufficiency proof of the Viability Theorem and again in the proof that the viability kernel is a nested intersection of finite-horizon sets, in both cases used to establish that a sequence of admissible trajectories is uniformly bounded on any fixed interval $[0, M]$, which is what licenses the Arzelà–Ascoli step in each proof. The bound here depends only on $\|\gamma(0)\|$, c , and t , uniformly across every trajectory satisfying the same growth hypothesis, exactly the uniformity those proofs needed and asserted without deriving.*

Lemma C.2 (Discrete Grönwall inequality). *Let $a_0, a_1, a_2, \dots \geq 0$ satisfy $a_{k+1} \leq (1 + hc) a_k + hc$ for constants $h, c > 0$. Then $a_k \leq (1 + a_0)(1 + hc)^k - 1$ for all $k \geq 0$.*

Proof. By induction. The case $k = 0$ is immediate. Assume $a_k \leq (1 + a_0)(1 + hc)^k - 1$. Then $a_{k+1} \leq (1 + hc)a_k + hc \leq (1 + hc)[(1 + a_0)(1 + hc)^k - 1] + hc = (1 + a_0)(1 + hc)^{k+1} - (1 + hc) + hc$, and $-(1 + hc) + hc = -1$, giving $a_{k+1} \leq (1 + a_0)(1 + hc)^{k+1} - 1$. \square

Remark C.2. *This is the discrete analogue used in the Euler polygon construction of the same sufficiency proof, where the piecewise-linear approximate trajectories x_0, x_1, x_2, \dots satisfy exactly this kind of recursive bound once the growth hypothesis is applied at each step, and $(1 + hc)^k \rightarrow e^{ckh}$ as $h \rightarrow 0$ with kh fixed recovers the continuous bound of Proposition C.1 in the limit, which is what allows the discrete and continuous versions of the argument to agree in the passage to the limit $h \rightarrow 0$.*

References

Bibliography

- [1] J.-P. Aubin, *Dynamic Core of Fuzzy Dynamical Cooperative Games*, Annals of Dynamic Games, 2001.
- [2] J.-P. Aubin, *Viability Theory*, Birkhäuser, 1991.
- [3] J.-P. Aubin, A. Cellina, *Differential Inclusions*, Springer-Verlag, 1984.
- [4] J.-P. Aubin, H. Frankowska, *Set-Valued Analysis*, Birkhäuser, 1990.
- [5] C. Alexander, S. Ishikawa, M. Silverstein, *A Pattern Language*, Oxford University Press, 1977.
- [6] R. Bellman, *Dynamic Programming*, Princeton University Press, 1957.
- [7] A. F. Filippov, *Differential Equations with Discontinuous Righthand Sides*, Kluwer Academic Publishers, 1988.
- [8] A. F. Filippov, "On certain questions in the theory of optimal control," *SIAM Journal on Control*, 1 (1962).
- [9] T. H. Grönwall, "Note on the derivatives with respect to a parameter of the solutions of a system of differential equations," *Annals of Mathematics*, 20 (1919).
- [10] J. Holland, *Hidden Order: How Adaptation Builds Complexity*, Addison-Wesley, 1995.
- [11] S. Kauffman, *Investigations*, Oxford University Press, 2000.
- [12] K. Kuratowski, C. Ryll-Nardzewski, "A general theorem on selectors," *Bulletin de l'Académie Polonaise des Sciences*, 13 (1965).
- [13] M. Levin, "Bioelectric signaling: Reprogrammable circuits underlying embryogenesis, regeneration, and cancer," *Cell*, 184(8), 2021.
- [14] A. M. Lyapunov, *The General Problem of the Stability of Motion*, 1892 (trans. Taylor & Francis, 1992).
- [15] R. Paley, N. Wiener, A. Zygmund, "Notes on random functions," *Mathematische Zeitschrift*, 37 (1933).
- [16] J. Pearl, *Causality: Models, Reasoning, and Inference*, Cambridge University Press, 2000.

- [17] H. Poincaré, *Sur le problème des trois corps et les équations de la dynamique*, Acta Mathematica, 13 (1890).
- [18] S. Shreve, *Stochastic Calculus for Finance II: Continuous-Time Models*, Springer, 2004.
- [19] V. Utkin, *Sliding Modes in Control and Optimization*, Springer-Verlag, 1992.