Continuous Capacities on Continuous State Spaces*

Jean Goubault-Larrecq

LSV, ENS Cachan, CNRS, INRIA Futurs
61, avenue du président-Wilson, F-94235 Cachan goubault@lsv.ens-cachan.fr

Abstract. We propose axiomatizing some stochastic games, in a continuous state space setting, using continuous belief functions, resp. plausibilities, instead of measures. Then, stochastic games are just variations on continuous Markov chains. We argue that drawing at random along a belief function is the same as letting the probabilistic player P play first, then letting the non-deterministic player C play demonically. The same holds for an angelic C, using plausibilities instead. We then define a simple modal logic, and characterize simulation in terms of formulae of this logic. Finally, we show that (discounted) payoffs are defined and unique, where in the demonic case, P maximizes payoff, while C minimizes it.

1 Introduction

Consider Markov chains: these are transition systems, which evolve from state $x \in X$ by drawing the next state y in the state space X according to some probability distribution $\theta(x)$. One may enrich this model to take into account decisions made by a *player* P, which can take actions ℓ in some set L. In state $x \in X$, P chooses an action $\ell \in L$, and draws the next state y according to a probability distribution $\theta_\ell(x)$ depending on $\ell \in L$: these are *labeled Markov processes* (LMPs) [8]. Adding rewards $r_\ell(x)$ on taking action ℓ from state x yields *Markov decision processes* [11]. The main topic there is to evaluate strategies that maximize the expected payoff, possibly discounted.

These notions have been generalized in many directions. Consider stochastic games, where there is not one but several players, with different goals. In security protocols, notably, it is meaningful to assume that the honest agents collectively define a player P as above, who may play probabilistically, and that attackers define a second player C, who plays *non-deterministically*. Instead of drawing the next state at random, C deliberately chooses its next state, typically to minimize P's expected payoff or to maximize the probability that a bad state is reached—this is *demonic* non-determinism.

A nice idea of F. Laviolette and J. Desharnais (private comm., 2003), which we develop, is that the theory of these games could be simplified by relaxing the requirements of Markov chains: if $\nu = \theta_\ell(x)$ is not required to be a measure, but the additivity requirement is relaxed to sub-additivity (i.e., $\nu(A)+\nu(B) \leq \nu(A \cup B)$ for disjoint measurable sets A, B), then such "preprobabilities" include both ordinary probabilities and the following funny-looking *unanimity game* $\mathfrak{u}_A$, which represents the demonic non-deterministic choice of an element from the set A: the preprobability $\mathfrak{u}_A(B)$ of drawing an element in B is 1 if $A \subseteq B$, 0 otherwise. The intuition is as follows. Assume that,

* Partially supported by the INRIA ARC ProNoBis.

starting from state x, you would like the next state y to be in B. A demonic adversary C will then strive to pick y outside B. Now if C's moves are given by $\delta_\ell(x) = \mathfrak{u}_A$, then either $A \not\subseteq B$, then it is C's interest to pick y from $A \setminus B$, so that the preprobability that y be in B is 0; or $A \subseteq B$, then C is forced to play $y \in B$, and the preprobability is 1.

However, sub-additive set functions are not quite the right notion; and second (which does not detract from F. Laviolette and J. Desharnais' great intuition), the right notions had been invented by economists in the 1950s under the name of "cooperative game with transferable utility" [22] and by statisticians in the 1960s under the names of belief functions and plausibilities, while capacities and Choquet integration are even more ancient [4]. A nice survey is [13]. These notions are well-known in discrete state spaces. Our generalization to topological spaces is new, and non-trivial. The spaces we consider include finite spaces as well as infinite ones such as $\mathbb{R}^n$, but also cpos and domains.

Outline. We introduce necessary mathematical notions in Section 2. We then develop the theory of continuous games, and continuous belief functions in particular in Section 3, showing in a precise sense how the latter model both probabilistic and demonic non-deterministic choice. We then recall the Choquet integral in Section 4, and show how taking averages reflects the fact that C aims at minimizing P's gains. We briefly touch the dual notion of plausibilities (angelic non-determinism) in passing. Finally, we define ludic transition systems, the analogue of Markov chains, except using continuous games, in Section 5, and define a notion of simulation topologies. We show that the coarsest simulation topology is exactly that defined by a simple modal logic, à la Larsen-Skou [19]. This illustrates how continuous games allow us to think of certain stochastic games as really being just LMPs, only with a relaxed notion of probability.

This work is a summary of most of Chapters 1-9 of [14], in which all proofs, and many more results can be found.

Related Work. Many models of Markov chains or processes, or stochastic games are discrete or even finite-state. Desharnais *et al.* [8] consider LMPs over *analytic* spaces, a class of topological spaces that includes not only finite spaces but also spaces such as $\mathbb{R}^n$. They show an extension of Larsen and Skou's Theorem [19]: two states are probabilistically bisimilar iff they satisfy the same formulae of the logic whose formulae are $F ::= \top | F \wedge F | [\ell]_{>r} F$, where $[\ell]_{>r} F$ is true at those states x where the probability $\theta_\ell(x)(\llbracket F \rrbracket_\theta)$ of going to some state satisfying F by doing action ℓ is greater than r. This is extended to any measurable space through *event bisimulations* in [5].

Mixing probabilistic (player P) and non-deterministic (C) behavior has also received some attention. This is notably at the heart of the *probabilistic I/O automata* of Segala and Lynch [25]. The latter can be seen as labeled Markov processes with discrete probability distributions $\theta_\ell(x)$ (i.e., linear combinations of Dirac masses), where the set L of actions is partitioned into internal (hidden) actions and external actions. While P controls the latter, the former represent non-deterministic transitions, i.e., under the control of C. Our model of stochastic games is closer to the strictly alternating variant of probabilistic automata, where at each state, a non-deterministic choice is made among several distributions, then the next state is drawn at random according to the chosen distribution. I.e., C plays, then P, and there is no intermediate state where C would have played but not P. This is similar to the model by Mislove *et al.* [21], who consider state spaces that are continuous cpos. In our model, this is the other way around: in each state,

P draws at random a possible choice set for C, who then picks non-deterministically from it. Additionally, our model accommodates state spaces that are discrete, or continuous cpos, or topological spaces such as $\mathbb{R}^n$, without any change to be made. Mislove *et al.* [21] consider a model where non-determinism is chaotic, i.e., based on a variant of Plotkin's powerdomain. We concentrate on demonic non-determinism, which is based on the Smyth powerdomain instead. For angelic non-determinism, see [14, chapitre 6], and [14, chapitre 7] for chaotic non-determinism.

Bisimulations have been studied in the above models. There are many variants on probabilistic automata [26, 16, 23]. Mislove *et al.* [21] show that (bi)simulation in their model is characterized by a logic similar to [8], with an added disjunction operator. Our result is similar, for a smaller logic, with one less modality. Segala and Turrini [27] compare various notions of bisimulations in these contexts.

We have already mentioned cooperative games and belief functions. See the abundant literature [6, 7, 28, 13, 24, 2]. We view belief functions as generalized probabilities; the competing view as a basis for a theory of evidence is incompatible [15].

An obvious approach to studying probabilistic phenomena is to turn to measure theory and measurable spaces, see e.g. [3]. However, we hope to demonstrate that the theory of cooperative games in the case of infinite state spaces X is considerably more comfortable when X is a topological space, and we only measure opens instead of Borel subsets. This is in line with the theory of continuous valuations [17], which has had considerable success in semantics.

We use Choquet integration to integrate along capacities ν [4]. This is exactly the notion that Tix [30] used more recently, too, and coincides with the Jones integral [17] for integration along continuous valuations. Finally, we should also note that V. Danos and M. Escardo have also come up (private comm.) with a notion of integration that generalizes Choquet integration, at least when integrating with respect to a convex game.

Acknowledgments. Thanks to F. Laviolette, J. Desharnais, V. Danos, P. Panangaden, Ph. Scott, M. Escardo, and the many others who expressed their support. Thanks to the anonymous referees for their helpful comments.

2 Preliminaries

Our state spaces X are topological spaces. We assume the reader to be familiar with (point-set) topology, in particular topology of T_0 but not necessarily Hausdorff spaces. See [12, 1, 20] for background. Let $int(A)$ denote the interior of A, $cl(A)$ its closure.

The *Scott topology* on a poset X, with ordering $\leq$, has as opens the upward-closed subsets U (i.e., $x \in U$ and $x \leq y$ imply $y \in U$) such that for every directed family $(x_i)_{i \in I}$ having a least upper bound $\sup_{i \in I} x_i$ inside U, some x_i is already in U. The *way-below* relation $\ll$ is defined by $x \ll y$ iff for any directed family $(z_i)_{i \in I}$ with a least upper bound z such that $y \leq z$, then $x \leq z_i$ for some $i \in I$. A poset is *continuous* iff $\Downarrow y = \{x \in X | x \ll y\}$ is directed, and has x as least upper bound. Then every open U can be written $\bigcup_{x \in U} \Uparrow x$, where $\Uparrow x = \{y \in X | x \ll y\}$.

Every topological space X has a specialization quasi-ordering $\leq$, defined by: $x \leq y$ iff every open that contains x contains y. X is T_0 iff $\leq$ is a (partial) ordering. That of the Scott topology of a quasi-ordering $\leq$ is $\leq$ itself. A subset $A \subseteq X$ is *saturated* iff A

is the intersection of all opens that contain it; alternatively, iff A is upward-closed in $\leq$. Every open is upward-closed. Let $\uparrow A$ denote the upward-closure of A under a quasi-ordering $\leq$, $\downarrow A$ its downward-closure. A T_0 space is *sober* iff every irreducible closed subset is the closure $cl\{x\} = \downarrow x$ of a (unique) point x. The Hofmann-Mislove Theorem implies that every sober space is *well-filtered* [18], i.e., given any filtered family of saturated compacts $(Q_i)_{i \in I}$ in X, and any open U, $\bigcap_{i \in I} Q_i \subseteq U$ iff $Q_i \subseteq U$ for some $i \in I$. In particular, $\bigcap_{i \in I} Q_i$ is saturated compact. X is *locally compact* iff whenever $x \in U$ (U open) there is a saturated compact Q such that $x \in int(Q) \subseteq Q \subseteq U$. Every continuous cpo is sober and locally compact in its Scott topology. We shall consider the space $\mathbb{R}$ of all reals with the Scott topology of its natural ordering $\leq$. Its opens are $\emptyset$, $\mathbb{R}$, and the intervals $(t, +\infty)$, $t \in \mathbb{R}$. $\mathbb{R}$ is a stably locally compact, continuous cpo. Since we equip $\mathbb{R}$ with the Scott topology, our *continuous* functions $f : X \to \mathbb{R}$ are those usually called *lower semi-continuous* in the mathematical literature.

We call *capacity* on X any function ν from $\mathcal{O}(X)$, the set of all opens of X, to $\mathbb{R}^+$, such that $\nu(\emptyset) = 0$ (a.k.a., a *set function*.) A *game* ν is a *monotonic capacity*, i.e., $U \subseteq V$ implies $\nu(U) \leq \nu(V)$. (The name "game" is unfortunate, as there is no obvious relationship between this and games as they are usually defined in computer science, in particular with stochastic games. The name stems from cooperative games in economics, where X is the set of players, not states.) A *valuation* is a *modular* game ν, i.e., one such that $\nu(U \cup V) + \nu(U \cap V) = \nu(U) \cap \nu(V)$ for every opens U, V. A game is *continuous* iff $\nu(\bigcup_{i \in I} U_i) = \sup_{i \in I} \nu(U_i)$ for every directed family $(U_i)_{i \in I}$ of opens. Continuous valuations have a convenient theory that fits topology well [17, 18].

The *Dirac valuation* δ_x at $x \in X$ is the continuous valuation mapping each open U to 1 if $x \in U$, to 0 otherwise. (Note that $\delta_x = \mathfrak{u}_{\{x\}}$, by the way.) A finite linear combination $\sum_{i=1}^n a_i \delta_{x_i}$, $a_i \in \mathbb{R}^+$, is a *simple valuation*. All simple valuations are continuous. Conversely, Jones' Theorem [17, Theorem 5.2] states that, if X is a continuous cpo, then every continuous valuation ν is the least upper bound $\sup_{i \in I} \nu_i$ of a directed family $(\nu_i)_{i \in I}$ of simple valuations way-below ν. Continuous valuations are canonically ordered by $\nu \leq \nu'$ iff $\nu(U) \leq \nu'(U)$ for every open U of X.

3 Continuous Games, and Belief Functions

Defining the "preprobabilities" alluded to in the introduction is best done by strengthening super-additivity. A game ν on X on X is *convex* iff $\nu(U \cup V) + \nu(U \cap V) \geq \nu(U) + \nu(V)$ for every opens U, V. It is *concave* if the opposite inequality holds. Convex games are a cornerstone of economic theory. E.g., Shapley's Theorem states that (on a finite space) the core $\{p \text{ valuation on } X | \nu \leq p, \nu(X) = p(X)\}$ of any convex game ν is non-empty, which implies the existence of economic equilibria [13, 22]. But this has only been studied on discrete spaces (finiteness is implicit in [13], notably). Finite, and more generally discrete spaces are sets X, equipped with the discrete topology, so one may see our topological approach as a generalization of previous approaches.

Recall that the unanimity game $\mathfrak{u}_A$ is defined by $\mathfrak{u}_A(U) = 1$ if $A \subseteq U$, $\mathfrak{u}_A(U) = 0$ otherwise. Clearly, $\mathfrak{u}_A$ is convex. It is in fact more. Call a game ν *totally convex* (the standard name, when X is discrete, i.e., when U_i is an arbitrary subset of X, is "totally monotonic"; we changed the name so as to name total concavity the dual of total monotonicity) iff:

$$\nu\left(\bigcup_{i=1}^{n} U_i\right) \geq \sum_{I\subseteq\{1,\ldots,n\}, I\neq\emptyset} (-1)^{|I|+1}\nu\left(\bigcap_{i\in I} U_i\right) \qquad (1)$$

for every finite family $(U_i)_{i=1}^n$ of opens ($n \geq 1$), where $|I|$ denotes the cardinality of I. A *belief function* is a totally convex game. The dual notion of *total concavity* is obtained by replacing $\bigcup$ by $\bigcap$ and conversely in (1), and turning $\geq$ into $\leq$. A *plausibility* is a totally concave game. If $\geq$ is replaced by $=$ in (1), then we retrieve the familiar inclusion-exclusion principle from statistics. In particular any (continuous) valuation is a (continuous) belief function. Clearly, any belief function is a convex game. The converses of both statements fail: On $X = \{1, 2, 3\}$ with the discrete topology, $\mathfrak{u}_{\{1,2\}}$ is a belief function but not a valuation, and $\frac{1}{2}(\mathfrak{u}_{\{1,2\}} + \mathfrak{u}_{\{1,3\}} + \mathfrak{u}_{\{2,3\}} - \mathfrak{u}_{\{1,2,3\}})$ is a convex game but not a belief function.

When X is finite, it is well-known [13] that any capacity ν can be written $\sum_{A\neq\emptyset, A\subseteq X} \alpha_A \mathfrak{u}_A$ for some coefficients $\alpha_A \in \mathbb{R}$, in a unique way. Also, ν is a belief function iff all coefficients are non-negative. An interpretation of this formula is that ν is essentially a probabilistic choice of some non-empty subset A, with probability α_A, from which C can choose an element $y \in A$ non-deterministically.

Our first result is to show that this result transfers, in some form, to the general topological case. Let $\mathcal{Q}(X)$ be the *Smyth powerdomain* of X, i.e., the space of all non-empty compact saturated subsets Q of X, ordered by reverse inclusion $\supseteq$. $\mathcal{Q}(X)$ is equipped with its Scott topology, and is known to provide an adequate model of demonic non-determinism in semantics [1]. When X is well-filtered and locally compact, $\mathcal{Q}(X)$ is a continuous cpo. Its Scott topology is generated by the basic open sets $\Box U = \{Q \in \mathcal{Q}(X) | Q \subseteq U\}$, U open in X.

The relevance of $\mathcal{Q}(X)$ here can be obtained by realizing that a finite linear combination $\sum_{i=1}^n a_i \mathfrak{u}_{A_i}$ with positive coefficients is a *continuous* belief function iff every subset A_i is compact; and that $\mathfrak{u}_{A_i} = \mathfrak{u}_{\uparrow A_i}$. Any such linear combination that is continuous is therefore of the form $\sum_{i=1}^n a_i \mathfrak{u}_{Q_i}$, with $Q_i \in \mathcal{Q}(X)$. We call such belief functions *simple*. Returning to the interpretation above, this can be intuitively seen as a probabilistic choice of some set Q_i with probability a_i, from which C will choose $y \in Q_i$; additionally, Q_i is an element of $\mathcal{Q}(X)$, the traditional domain for *demonic* non-determinism.

So any simple belief function ν can be matched with a (simple) valuation $\nu^* = \sum_{i=1}^n a_i \delta_{Q_i}$ on $\mathcal{Q}(X)$. Note that $\nu^*(\Box U) = \nu(U)$ for every open U of X. This is exactly the sense in which continuous belief functions are essentially continuous valuations on the space $\mathcal{Q}(X)$ of non-deterministic choices.

Theorem 1. *For any continuous valuation P on $\mathcal{Q}(X)$, the capacity ν defined by $\nu(U) = P(\Box U)$ is a continuous belief function on X.*

Conversely, let X be a well-filtered and locally compact space. For every continuous belief function ν on X there is a unique continuous valuation ν^* on $\mathcal{Q}(X)$ such that $\nu(U) = \nu^*(\Box U)$ for every open U of X.

Proof. (Sketch.) The first part follows by computation. For the second part, observe that $\bigcup_{i=1}^n \Box U_i \subseteq \bigcup_{j=1}^m \Box V_j$ iff for every i, $1 \leq i \leq n$, there exists j, $1 \leq j \leq m$, such that

$U_i \subseteq V_j$. Thus, the function P given by $P\left(\bigcup_{i=1}^n \Box U_i\right) = \sum_{I \subseteq \{1,\ldots,n\}, I \neq \emptyset} (-1)^{|I|+1}$ $\nu\left(\bigcap_{i \in I} U_i\right)$ is well-defined and monotonic. Let $\nu^*(\mathcal{U})$ be the least upper bound of $P\left(\bigcup_{Q \in \mathcal{J}} \Box int(Q)\right)$, when $\mathcal{J}$ ranges over finite subsets of $\mathcal{U}$: ν^* is monotonic, continuous, $\nu^*(\Box U) = P(\Box U) = \nu(U)$, and fairly heavy computation shows that ν^* is modular. Uniqueness is easy. $\Box$

Next, we show that this bijection is actually an isomorphism, i.e., it also preserves order and therefore the Scott topology. To this end, define the ordering $\leq$ on all capacities, not just valuations, by $\nu \leq \nu'$ iff $\nu(U) \leq \nu'(U)$ for every open U of X. We start by characterizing it in the manner of Jones' splitting lemma. This [17, Theorem 4.10] states that $\sum_{i=1}^m a_i \delta_{x_i} \leq \sum_{j=1}^n b_j \delta_{y_j}$ iff there is matrix $(t_{ij})_{\substack{1 \leq i \leq m \\ 1 \leq j \leq n}}$ of coefficients in $\mathbb{R}^+$ such that $\sum_{j=1}^n t_{ij} = a_i$ for each i, $\sum_{i=1}^m t_{ij} \leq b_j$ for each j, and whenever $t_{ij} \neq 0$ then $x_i \leq y_j$. (Jones proves it for cpos, but it holds on any topological space [29, Theorem 2.4, Corollary 2.6].) We show:

Lemma 1 (Splitting Lemma). *$\sum_{i=1}^m a_i \mathfrak{u}_{Q_i} \leq \sum_{j=1}^n b_j \mathfrak{u}_{Q'_j}$ iff there is matrix $(t_{ij})_{\substack{1 \leq i \leq m \\ 1 \leq j \leq n}}$ of coefficients in $\mathbb{R}^+$ such that $\sum_{j=1}^n t_{ij} = a_i$ for each i, $\sum_{i=1}^m t_{ij} \leq b_j$ for each j, and whenever $t_{ij} \neq 0$ then $Q_i \supseteq Q'_j$.*

It follows that: (A) for any two simple belief functions ν, ν' on X, $\nu \leq \nu'$ iff $\nu^* \leq \nu'^*$, since the two are equivalent to the existence of a matrix $(t_{ij})_{\substack{1 \leq i \leq m \\ 1 \leq j \leq n}}$ satisfying the same conditions. This can be extended to all continuous belief functions, see below. Let $\mathbf{Cd}_{\leq 1}(X)$ be the space of continuous belief functions ν on X with $\nu(X) \leq 1$, ordered by $\leq$. Let $\mathbf{V}_{\leq 1}(X)$ the subspace of continuous valuations. We have:

Theorem 2. *Let X be well-filtered and locally compact. Every continuous belief function ν on X is the least upper bound of a directed family of simple belief functions ν_i way-below ν. $\mathbf{Cd}_{\leq 1}(X)$ is a continuous cpo.*

It follows that continuous belief functions are really the same thing as (sub-)probabilities over the set of demonic choice sets $Q \in \mathcal{Q}(X)$.

Theorem 3. *Let X be well-filtered and locally compact. The function $\nu \mapsto \nu^*$ defines an order-isomorphism from $\mathbf{Cd}_{\leq 1}(X)$ to $\mathbf{V}_{\leq 1}(\mathcal{Q}(X))$.*

As a side note, (up to the ≤ 1 subscript) $\mathbf{V}_{\leq 1}(\mathcal{Q}(X))$ is exactly the space into which Edalat [10] embeds a space of measures on X. The above Theorem states that the space of objects for which we can do this is exactly $\mathbf{Cd}_{\leq 1}(X)$.

Dually, we may mix probabilistic choice with angelic non-determinism. Space does not permit us to describe this in detail, see [14, chapitre 6]. The point is that the space $\mathbf{Pb}_{\leq 1}(X)$ of continuous plausibilities is order-isomorphic to $\mathbf{V}_{\leq 1}(\mathcal{H}_u(X))$, whenever X is stably locally compact, where the (topological) *Hoare powerdomain* $\mathcal{H}_u(X)$ of X is the set of non-empty closed subsets of X, with the upper topology of the inclusion ordering, generated by the subbasic sets $\Diamond U = \{F \in \mathcal{H}(X) | F \cap U \neq \emptyset\}$, U open in X. The argument goes through a nice notion of convex-concave duality, which intuitively exchanges good (concave) and evil (convex). The case of chaotic non-determinism is more complex, see [14, chapitre 7].

4 Choquet Integration

We introduce the standard notion of integration along games ν. This is mostly well-known [13]; adapting to the topological case is easy, so we omit proofs [14, chapitre 4].

Let ν be a game on X, and f be continuous from X to $\mathbb{R}$. Recall that we equip $\mathbb{R}$ with its Scott topology, so that f is really what is known otherwise as *lower semi-continuous*. Assume f bounded, too, i.e., $\inf_{x \in X} f(x) > -\infty$, $\sup_{x \in X} f(x) < +\infty$. The *Choquet integral* of f along ν is:

$$\oint^{C}_{x \in X} f(x) d\nu = \int_{0}^{+\infty} \nu(f^{-1}(t, +\infty)) dt + \int_{-\infty}^{0} [\nu(f^{-1}(t, +\infty)) - \nu(X)] dt \quad (2)$$

where both integrals on the right are improper Riemann integrals. This is well-defined, since $f^{-1}(t, +\infty)$ is open for every $t \in \mathbb{R}$ by assumption, and ν measures opens. Also, since f is bounded, the improper integrals above really are ordinary Riemann integrals over some closed intervals. The function $t \mapsto \nu(f^{-1}(t, +\infty))$ is decreasing, and every decreasing (even non-continuous, in the usual sense) function is Riemann-integrable, therefore the definition makes sense.

An alternate definition consists in observing that any *step function* $\sum_{i=0}^{n} a_i \chi_{U_i}$, where $a_0 \in \mathbb{R}$, $a_1, \ldots, a_n \in \mathbb{R}^+$, $X = U_0 \supseteq U_1 \supseteq \ldots \supseteq U_n$ is a decreasing sequence of opens, and χ_U is the indicator function of U ($\chi_U(x) = 1$ if $x \in X$, $\chi_U(x) = 0$ otherwise) is continuous, and of integral along ν equal to $\sum_{i=0}^{n} a_i \nu(U_i)$—for *any* game ν. It is well-known that every bounded continuous function f can be written as the least upper bound of a sequence of step functions $f_K = a + \frac{1}{2^K} \sum_{k=1}^{\lfloor (b-a)2^K \rfloor} \chi_{f^{-1}(a + \frac{k}{2^K}, +\infty)}(x)$, $K \in \mathbb{N}$, where $a = \inf_{x \in X} f(x)$, $b = \sup_{x \in X} f(x)$. Then the integral of f along ν is the least upper bound of the increasing sequence of the integrals of f_K along ν.

The main properties of Choquet integration are as follows. First, the integral is increasing in its function argument: if $f \leq g$ then the integral of f along ν is less than or equal to that of g along ν. If ν is continuous, then integration is also Scott-continuous in its function argument. The integral is also monotonic and Scott-continuous in the game ν, provided the function we integrate takes only non-negative values, or provided ν is *normalized*, i.e., $\nu(X) = 1$. Integration is linear in the game, too, so integrating along $\sum_{i=1}^{n} a_i \nu_i$ is the same as taking the integrals along each ν_i, and computing the obvious linear combination. However, Choquet integration is *not* linear in the function integrated, unless the game ν is a valuation. Still, it is *positively homogeneous*: integrating αf for $\alpha \in \mathbb{R}^+$ yields α times the integral of f. It is additive on *comonotonic* functions $f, g : X \to \mathbb{R}$ (i.e., there is no pair $x, x' \in X$ such that $f(x) < f(x')$ and $g(x) > g(x')$). It is super-additive (the integral of $f + g$ is at least that of f plus that of g) when ν is convex, in particular when ν is a belief function, and sub-additive when ν is concave. See [13] for the finite case, [14, chapitre 4] for the topological case.

One of the most interesting things is that integrating with respect to a unanimity game consists in taking minima. This suggests that unanimity games indeed model some *demonic* form of non-determinism. Imagine $f(x)$ is the amount of money you gain by going to state x. The following says that taking the average amount of money with respect to a demonic adversary C will give you back the least amount possible.

Proposition 1. *For any continuous $f : X \to \mathbb{R}^+$,*

$$\oint_{x \in X} f(x) d\mathfrak{u}_A = \inf_{x \in A} f(x)$$

Moreover, if A is compact, then the inf is attained: this equals $\min_{x \in A} f(x)$.

Since Choquet integration is linear in the game, the integral of f along a simple belief function $\sum_{i=1}^n a_i \mathfrak{u}_{Q_i}$ yields $\sum_{i=1}^n a_i \min_{x \in Q_i} f(x)$: this is the expected min-value of f obtained by drawing Q_i at random with probability a_i (P plays) then letting C non-deterministically move to the state $x \in Q_i$ that minimizes the gain. We can generalize this to non-discrete probabilities over $\mathcal{Q}(X)$ by using the $\nu \mapsto \nu^*$ isomorphism:

Theorem 4. *For any bounded continuous function $f : X \to \mathbb{R}$, let f_* be the function from $\mathcal{Q}(X)$ to $\mathbb{R}$ defined by $f_*(Q) = \min_{x \in Q} f(x)$. Say that a capacity ν is* linearly extensible from below *if and only if there is continuous valuation P on $\mathcal{Q}(X)$ with:*

$$\oint_{x \in X} f(x) d\nu = \oint_{Q \in \mathcal{Q}(X)} f_*(Q) dP \tag{3}$$

for every bounded continuous f. If X is well-filtered and locally compact, then the capacities that are linearly extensible from below are exactly the continuous belief functions, and P must be ν^* in (3).

It follows in particular that whenever ν is the least upper bound of a directed family $(\nu_i)_{i \in I}$ of simple belief functions ν_i, then integrating $f : X \to \mathbb{R}$ with respect to ν can be computed by taking least upper bounds of linear combinations of mins. Therefore the Choquet integral along continuous belief functions coincides with Edalat's *lower R-integral* [10], which was only defined for measures.

This can be dualized to the case of plausibilities ν, assuming X stably locally compact [14, théorème 6.3.17]. Then we talk about capacities that are linearly extensible from *above*. There is an isomorphism $\nu \mapsto \nu_*$ such that $\nu_*(\Diamond U) = \nu(U)$ for all U, and integrating f along ν amounts to integrating f^* along ν_*, where for every $F \in \mathcal{H}_u(X)$, $f^*(F) = \sup_{x \in F} f(x)$. (I.e., C now *maximizes* our gain.) Then the Choquet integral along continuous plausibilities coincides with Edalat's *upper R-integral* [10].

5 Ludic Transition Systems, Logic, Simulation, Rewards

Let $\mathbf{J}_{\leq 1}(X)$ be the space of all continuous games ν on X with $\nu(X) \leq 1$. This is equipped with its Scott topology. It will be practical to consider another topology. The *weak topology* on a subspace Y of $\mathbf{J}_{\leq 1}(X)$ is the topology generated by the subbasic open sets $[U > r] = \{\nu \in Y | \nu(U) > r\}$, U open in X, $r \in \mathbb{R}$. It is in general coarser than the Scott topology, and coincides with it when $Y = \mathbf{V}_{\leq 1}(X)$ and X is a continuous cpo [30, Satz 4.10]. One can show that the weak topology is exactly the coarsest that makes continuous all functionals mapping $\nu \in Y$ to the integral of f along ν, for all $f : X \to \mathbb{R}^+$ bounded continuous. (See [14, section 4.5] for details.)

By analogy with Markov kernels and LMPs, define a *ludic transition system* as a family $\theta = (\theta_\ell)_{\ell \in L}$, where L is a given set of *actions*, and each θ_ℓ is a continuous map from the state space X to $\mathbf{J}_{\leq 1\ wk}(X)$. (See [14, chapitres 8, 9] for missing details.) The main change is that, as announced in the introduction, we replace probability distributions by continuous games. One may object that LMPs are defined as *measurable*, not continuous, so that this definition overly restricts the class of transition systems we are considering. However, the mathematics are considerably cleaner when assuming continuity. Moreover, the weak topology is so weak that, for example, it only restrains θ_ℓ so that $x \mapsto \theta_\ell(x)(U)$ is continuous as a function from X to $\mathbb{R}^+$, equipped with its *Scott* topology; this certainly allows it to have jumps. Finally, one may argue, following Edalat [10], that any second countable locally compact Hausdorff space X can be embedded as a set of maximal elements of a continuous cpo (namely $\mathcal{Q}(X)$; other choices are possible) so that any measure on X extends to a continuous valuation on $\mathcal{Q}(X)$. This provides a theory of approximation of integration on X through domain theory. One may hope a similar phenomenon will apply to games—for some notion of games yet to be defined on Borel subsets, not opens.

Logic. Following [8, 5], define the logic $\mathcal{L}^{\top\wedge\vee}_{\text{open}}$ by the grammar shown right, where $\ell \in L$, $r \in \mathbb{Q} \cap [0,1]$ in the last line. Compared to [8, 5], we only have one extra disjunction operator. The same logic, with disjunction, is shown to characterize simulation for LMPs in [9, Section 2.3].

$$
\begin{array}{rcll}
F & ::= & \top & \text{true} \\
 & \mid & F \wedge F & \text{conjunction (and)} \\
 & \mid & F \vee F & \text{disjunction (or)} \\
 & \mid & [\ell]_{>r} F & \text{modality}
\end{array}
$$

Let $[\![F]\!]_\theta$ be the set of states $x \in X$ where F holds: $[\![\top]\!]_\theta = X$, $[\![F_1 \wedge F_2]\!]_\theta = [\![F_1]\!]_\theta \wedge [\![F_2]\!]_\theta$, $[\![F_1 \vee F_2]\!]_\theta = [\![F_1]\!]_\theta \vee [\![F_2]\!]_\theta$, and $[\![[\ell]_{>r} F]\!]_\theta = \delta_\ell^{-1}[[\![F]\!]_\theta > r]$ is the set of states x such that the preprobability $\delta_\ell([\![F]\!]_\theta)$ that the next state y will satisfy F on firing an ℓ action is strictly greater than r. Note that this is well-defined, precisely because δ_ℓ is continuous from X to a space of games with the weak topology. Also, it is easy to see that $[\![F]\!]_\theta$ is always open.

Simulation. Now define *simulation* in the spirit of event bisimulation [5] (we shall see below why we do not call it *bi*simulation). For any topology $\mathcal{O}$ on X coarser than that of X, let $X : \mathcal{O}$ be X equipped with the topology $\mathcal{O}$. A *simulation topology* for θ is a topology $\mathcal{O}$ on X, coarser than that of X, such that δ_ℓ is continuous from $X : \mathcal{O}$ to $\mathbf{J}_{\leq 1\ wk}(X : \mathcal{O})$, i.e., $\delta_\ell^{-1}[U > r] \in \mathcal{O}$ for each $U \in \mathcal{O}$ and each $r \in \mathbb{R}$. (A close notion was introduced in [31, Theorem 29].) One non-explanation for this definition is to state that this is exactly event bisimulation [5], only replacing σ-algebras by topologies. A better explanation is to revert back to Larsen and Skou's original definition of probabilistic bisimulation in terms of an algebra of *tests* (in slightly more abstract form). A (bi)simulation should not be thought as an arbitrary equivalence relation, rather as one generated from a collection Tst of tests, which are subsets A of X: $x \in X$ passes the test iff $x \in A$, it fails it otherwise. Two elements are equivalent iff they pass the same tests. Now in a continuous setting it only makes sense that the tests be open: any open U defines a continuous predicate χ_U from X to the Sierpiński space $\mathbb{S} = \{0, 1\}$ (with the Scott topology of $0 \leq 1$), and conversely. Let $\mathcal{O}_{Tst}$ be the topology generated by the tests Tst. It is sensible to require that $\delta_\ell^{-1}[U > r]$ be a test, too, at least when U is a finite union of finite intersections of tests (for the general case, appeal to the fact that $\delta_\ell(x)$ is continuous, and that any open can be approximated by such a finite union): one can indeed test whether $x \in \delta_\ell^{-1}[U > r]$ by firing transitions according to the preprobability $\delta_\ell(x)$, and test (e.g., by sampling, knowing that if $\delta_\ell(x)$ is a belief function for example, then we are actually playing also against a demonic adversary C) whether our chances of getting to a state $y \in U$ exceed r. And this is essentially how we defined simulation topologies.

Every simulation topology $\mathcal{O}$ defines a specialization quasi-ordering $\preceq_{\mathcal{O}}$, which is the analogue of the standard notion of simulation here. (Note that in the case of event bisimulation, i.e., taking σ-algebras instead of topologies, $\preceq_{\mathcal{O}}$ would be an equivalence relation—because σ-algebras are closed under complements—justifying the fact that event bisimulation really is a bisimulation, while our notion is a simulation.) Write $\equiv_{\mathcal{O}} = \preceq_{\mathcal{O}} \cap \succeq_{\mathcal{O}}$ the equivalence associated with simulation $\preceq_{\mathcal{O}}$. Clearly, there is a coarsest (largest) simulation topology $\mathfrak{O}_\theta$. The following is then easy:

Theorem 5. *Let $\mathcal{O}$ be a simulation topology for θ on X. For any $F \in \mathcal{L}_{open}$, $[\![F]\!]_\theta \in \mathcal{O}$. In particular [Soundness], if $x \in [\![F]\!]_\theta$ and $x \preceq_{\mathcal{O}} y$ then $y \in [\![F]\!]_\theta$. Conversely [Completeness], the coarsest simulation topology $\mathfrak{O}_\theta$ is exactly that generated by the opens $[\![F]\!]_\theta$, $F \in \mathcal{L}^{\top\wedge\vee}_{open}$.*

This can be used, as is standard in the theory of Markov chains, to *lump* states. Given a topology $\mathcal{O}$, let $X/\mathcal{O}$ be the quotient space $X/{\equiv_{\mathcal{O}}}$, equipped with the finest topology such that $q_{\mathcal{O}} : X : \mathcal{O} \to X/\mathcal{O}$ is continuous. Let the *direct image* $f[\nu]$ of a game ν on X by a continuous map $f : X \to Y$ be $f[\nu](V) = \nu(f^{-1}(V))$. Taking direct images preserves monotonicity, modularity, (total) convexity, (total) concavity, and continuity.

Proposition 2. *Let $\mathcal{O}$ be a simulation topology for θ. The function $\theta_\ell/\mathcal{O}$ mapping $q_{\mathcal{O}}(x)$ to $q_{\mathcal{O}}[\theta_\ell(x)]$ is well defined and continuous from $X/\mathcal{O}$ to $\mathbf{J}_{\leq 1\ wk}(X/\mathcal{O})$ for every $\ell \in L$. The family $\theta/\mathcal{O} = (\theta_\ell/\mathcal{O})_{\ell \in L}$ is then a ludic transition system on $X/\mathcal{O}$, which we call the* lumped *ludic transition system.*

For any $F \in \mathcal{L}^{\top\wedge\vee}_{open}$ and $x \in X$, x and $q_{\mathcal{O}}(x)$ satisfy the same formulae: $q_{\mathcal{O}}([\![F]\!]_\theta) = [\![F]\!]_{\theta/\mathcal{O}}$, and $[\![F]\!]_\theta = q_{\mathcal{O}}^{-1}([\![F]\!]_{\theta/\mathcal{O}})$, in particular, $x \in [\![F]\!]_\theta$ iff $q_{\mathcal{O}}(x) \in [\![F]\!]_{\theta/\mathcal{O}}$.

Rewards and payoffs. A classical problem on Markov decision processes is to evaluate average payoffs. Since LMPs and ludic transition systems are so similar, we can do exactly the same. Imagine P plays according to a finite-state program Π, i.e., an automaton with *internal states* q, q' and transitions $q \stackrel{\ell}{\longrightarrow} q'$. Let $r_{q \stackrel{\ell}{\longrightarrow} q'} : X \to \mathbb{R}$ be a family of bounded continuous *reward* functions: we may think that $r_{q \stackrel{\ell}{\longrightarrow} q'}(x)$ is the amount of money P gains if she fires her internal transition $q \stackrel{\ell}{\longrightarrow} q'$, drawing the next state y at random along $\theta_\ell(x)$. Let $\gamma_{q \stackrel{\ell}{\longrightarrow} q'} \in (0, 1]$ be a family of so-called *discounts*. Define the average payoff, starting from state x when P is in its internal state q, by:

$$V_q(x) = \sup_{\ell, q' / q \stackrel{\ell}{\longrightarrow} q'} \left[r_{q \stackrel{\ell}{\longrightarrow} q'}(x) + \gamma_{q \stackrel{\ell}{\longrightarrow} q'} \oint_{y \in X} V_{q'}(y) d\theta_\ell(x) \right] \tag{4}$$

This formula would be standard if $\theta_\ell(x)$ were a probability distribution. What is less standard is what (4) means when $\theta_\ell(x)$ is a game. E.g., when $\theta_\ell(x)$ is a simple belief function $\sum_{i=1}^{n_\ell} a_{i\ell x} \mathfrak{u}_{Q_{i\ell x}}$, then:

$$V_q(x) = \sup_{\ell, q' / q \stackrel{\ell}{\longrightarrow} q'} \left[r_{q \stackrel{\ell}{\longrightarrow} q'}(x) + \gamma_{q \stackrel{\ell}{\longrightarrow} q'} \sum_{i=1}^{n_\ell} a_{i\ell x} \min_{y \in Q_{i\ell x}} V_{q'}(y) \right] \tag{5}$$

where we see that P has control over the visible transitions ℓ, and tries to maximize his payoff ($\sup$), while C will minimize it, and some averaging is taking place in-between. The equation (4) does not always have a solution in the family of all V_qs. But there are two cases where it has, similar to those encountered in Markov decision processes.

Theorem 6. *Assume θ is* standard*, i.e., $\theta_\ell(X)$ is always either 0 or 1, and the set $\{x \in X | \theta_\ell(x) = 0\}$ of* deadlock states *is open; or that $r_{q \xrightarrow{\ell} q'}(x) \geq 0$ for all q, ℓ, q', $x \in X$. Assume also that there are $a, b \in \mathbb{R}$ with $a \leq r_{q \xrightarrow{\ell} q'}(x), \gamma_{q \xrightarrow{\ell} q'} \leq b$ for all q, ℓ, q', $x \in X$. Then (4) has a unique solution in any of the following two cases:*
[Finite Horizon] If all paths in Π have bounded length.
[Discount] If there is a constant $\gamma \in (0, 1)$ such that $\gamma_{q \xrightarrow{\ell} q'} \leq \gamma$ for every q, ℓ, q'.

When θ_ℓ is a simple belief function, Equation (5) is then a Bellman-type equation that can be solved by dynamic programming techniques. Then observe that any continuous belief function is the directed lub of simple belief functions by Theorem 2, under mild assumptions. This offers a canonical way to approximate the average payoff V_q.

References

1. Samson Abramsky and Achim Jung. Domain theory. In S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum, editors, *Handbook of Logic in Computer Science*, volume 3, pages 1–168. Oxford University Press, 1994.
2. Gilbert W. Bassett Jr., Roger Koenker, and Gregory Kordas. Pessimistic portfolio allocation and Choquet expected utility. Available from `http://www.econ.uiuc.edu/~roger/research/risk/choquet.pdf`, January 2004.
3. Stefano Cattani, Roberto Segala, Marta Z. Kwiatkowska, and Gethin Norman. Stochastic transition systems for continuous state spaces and non-determinism. In V. Sassone, editor, *Proc. 8th Int. Conf. Foundations of Software Science and Computation Structures (FoSSaCS 2005)*. Springer-Verlag LNCS 3441, 125–139.
4. Gustave Choquet. Theory of capacities. *Annales de l'Institut Fourier*, 5:131–295, 1953–54.
5. Vincent Danos, Josée Desharnais, François Laviolette, and Prakash Panangaden. Bisimulation and cocongruence for probabilistic systems. *Information and Computation*, 204(4):503–523, April 2006. Special issue for selected papers from CMCS04, 22 pages.
6. A. P. Dempster. Upper and lower probabilities induced by a multivalued mapping. *Annals of Mathematical Statistics*, 38:325–339, 1967.
7. A. P. Dempster. A generalization of Bayesian inference. *Journal of the Royal Statistical Society*, B 30:205–247, 1968.
8. Josée Desharnais, Abbas Edalat, and Prakash Panangaden. Bisimulation for labelled Markov processes. *Information and Computation*, 179(2):163–193, 2002.
9. Josée Desharnais, Vineet Gupta, Radhakrishnan Jagadeesan, and Prakash Panangaden. Approximating labeled Markov processes. *Information and Computation*, 184(1):160–200, July 2003.
10. Abbas Edalat. Domain theory and integration. *Theoretical Computer Science*, 151:163–193, 1995.
11. Eugene A. Feinberg and Adam Schwartz, editors. *Handbook of Markov Decision Processes, Methods and Applications*. Kluwer, 2002. 565 pages.
12. Gerhard Gierz, Karl Heinrich Hofmann, Klaus Keimel, Jimmie D. Lawson, Michael Mislove, and Dana S. Scott. *A Compendium of Continuous Lattices*. Springer Verlag, 1980.

13. Itzhak Gilboa and David Schmeidler. Additive representation of non-additive measures and the Choquet integral. Discussion Papers 985, Northwestern University, Center for Mathematical Studies in Economics and Management Science, 1992.
14. Jean Goubault-Larrecq. Une introduction aux capacités, aux jeux et aux prévisions. `http://www.lsv.ens-cachan.fr/~goubault/ProNobis/pp.pdf`, January 2007. 516 pages.
15. Joseph Y. Halpern and Ronald Fagin. Two views of belief: Belief as generalized probability and belief as evidence. *Artificial Intelligence*, 54:275–317, 1992.
16. Hans A. Hansson and Bengt Jonsson. A calculus for communicating systems with time and probabilities. In *Proc. 11th IEEE Real-time Systems Symp.*, pages 278–287, Silver Spring, MD, 1990. IEEE Computer Society Press.
17. Claire Jones. *Probabilistic Non-Determinism*. PhD thesis, University of Edinburgh, 1990. Technical Report ECS-LFCS-90-105.
18. Achim Jung. Stably compact spaces and the probabilistic powerspace construction. In J. Desharnais and P. Panangaden, editors, *Domain-theoretic Methods in Probabilistic Processes*, volume 87 of *Electronic Lecture Notes in Computer Science*. Elsevier, 2004. 15pp.
19. Kim G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94:1–28, 1991.
20. Michael Mislove. Topology, domain theory and theoretical computer science. *Topology and Its Applications*, 89:3–59, 1998.
21. Michael Mislove, Joël Ouaknine, and James Worrell. Axioms for probability and nondeterminism. *Electronic Notes in Theoretical Computer Science*, 91(3):7–28, 2003. Proc. 10th Int. Workshop on Expressiveness in Concurrency (EXPRESS'03).
22. Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
23. A. Philippou, I. Lee, and O. Sokolsky. Weak bisimulation for probabilistic processes. In C. Palamidessi, editor, *Proc. CONCUR 2000*, pages 334–349. Springer-Verlag LNCS 1877, 2000.
24. David Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57:571–587, 1989.
25. Roberto Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. MIT Press, Cambridge, MA, 1996.
26. Roberto Segala and Nancy Lynch. Probabilistic simulations for probabilistic processes. *Nordic Journal of Computing*, 2(2):250–273, 1995.
27. Roberto Segala and Andrea Turrini. Comparative analysis of bisimulation relations on alternating and non-alternating probabilistic models. In *2nd Int. Conf. Quantitative Evaluaton of Systems (QEST 2005)*, pages 44–53, Torino, Italy, September 2005. IEEE Computer Society Press.
28. G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, NJ, USA, 1976.
29. Philipp Sünderhauf. Spaces of valuations as quasimetric domain. In A. Edalat, A. Jung, K. Keimel, and M. Kwiatkowska, editors, *Proceedings of the 3rd Workshop on Computation and Approximation (Comprox III)*, volume 13 of *Electronic Notes in Theoretical Computer Science*, Birmingham, England, September 1997. Elsevier.
30. Regina Tix. *Stetige Bewertungen auf topologischen Räumen*. Diplomarbeit, TH Darmstadt, June 1995.
31. Franck van Breugel, Michael Mislove, Joël Ouaknine, and James Worrell. Domain theory, testing and simulation for labelled markov processes. *Theoretical Computer Science*, 333(1-2):171–197, 2005.