

Demonic, angelic and unbounded probabilistic choices in sequential programs

AK McIver* and Carroll Morgan†

February 5, 2007

Abstract

Probabilistic predicate transformers extend standard predicate transformers by adding probabilistic choice to (transformers for) sequential programs; demonic nondeterminism is retained. For finite state spaces, the basic theory is set out elsewhere [17], together with a presentation of the probabilistic ‘healthiness conditions’ that generalise the ‘positive conjunctivity’ of ordinary predicate transformers.

Here we expand the earlier results beyond ordinary conjunctive transformers, investigating the structure of the transformer space more generally: as Back and von Wright [1] did for the standard (non-probabilistic) case, we nest deterministic, demonic and demonic/angelic transformers, showing how each subspace can be constructed from the one before. We show also that the results hold for infinite state spaces.

In the end we thus find characteristic healthiness conditions for the hierarchies of a system in which deterministic, demonic, probabilistic and angelic choices all coexist.

Keywords: Probability, predicate transformers, nondeterminism, verification, refinement, imperative programming, program derivation, Galois connection.

*McIver *anabel@comlab.ox.ac.uk* is a member of the Programming Research Group at the University of Oxford, and is supported by the EPSRC.

†Morgan *carroll@comlab.ox.ac.uk* was partially supported during this work by the Dutch *Specification and Transformation of Programs* (STOP) project.

1 Introduction

Probabilistic predicate transformers, or ‘expectation transformers’ as we call them here, were introduced by Kozen [6] as a logic for imperative probabilistic programming. Recently [17] we extended Kozen’s work to include demonic nondeterminism, thus placing it on a par with the general treatment of imperative nondeterminism in the ‘standard’ (non-probabilistic) case [2].

In this paper we make two contributions. First, we investigate systematically the structure of the whole expectation transformer space — not just the conjunctive-like part of it — revealing a hierarchy of program classes (Fig. 9 below), with the more complex being built from the simpler ones. (Back and von Wright [1] reveal a similar hierarchy for standard programs.) Further, we associate with each class of program a characteristic ‘healthiness condition’ in the expectation logic, showing it to be necessary and sufficient for membership of the class.

Second, we show that both our earlier work [17] and the current results extend to infinite state spaces, provided a continuity assumption is made.

Programs as standard predicate transformers can be presented as nested classes of increasing sophistication: the *deterministic* programs deliver only one final state for any initial state, and are characterised by disjunctivity and conjunctivity in the (Boolean) programming logic; *demonic* programs may contain ‘worst-case’ nondeterministic choices, and are characterised by conjunctivity alone; demonic/*angelic* programs may contain ‘best-case’ non-deterministic choice as well, and have been shown [1] to be characterised simply by monotonicity.

Probabilistic ‘expectation transformers’ can be similarly presented. We show here that deterministic (but probabilistic) programs have a predictable *distribution* of output behaviours; and their associated expectation transformers are characterised by ‘linearity’. Demonic programs are made by demonic choices between deterministic ones; their expectation transformers are characterised by ‘sublinearity’. And demonic/angelic programs are made by angelic choices between demonic ones, with their expectation transformers characterised by ‘semi-sublinearity’.

In Sec. 2 we review Kozen’s work and our recent extensions of it. Sections 3, 4 and 5 then build the classes of deterministic, demonic and demonic/angelic programs. And finally Sec. 6 makes the extension to infinite state spaces: in most cases the earlier results are unaffected, and we merely indicate how their proofs can be re-used here.

Throughout we use infix ‘.’ to denote function application and ‘:=’ for ‘is defined to be’. We use A, B to denote standard (Boolean) predicates, often punning them with the (characteristic) $\{0, 1\}$ -valued functions of the state space commensurate with the context and to avoid the syntactic clutter of an explicit embedding function. In particular for states $s, s' \dots$ we write $\{s, s' \dots\}$ for the set containing those states, for the predicate holding only for those states and for the function 1 on those states and 0 elsewhere. We use α, β for general expectations and S for the underlying state space. Other notation is introduced as necessary.

2 Review of existing results

2.1 Introduction

Because Kozen [6, 7] studied only deterministic probabilistic programs — those containing neither demonic nor angelic choice — his operational model for sequential probabilistic programs consists essentially of Markov processes [3]: a program’s execution from a given initial state is described by giving the resulting probability distribution over final states, and for example sequential composition is then a multiplication of appropriate Markov transition matrices.

Our extension [17] added demonic nondeterminism to that operational model in a simple way — programs are no longer Markov processes, but rather are a generalised form; and a computation evolves by first making a demonic choice, in effect selecting one distribution from a number of possibilities, and then using the selected distribution to resolve the immediately following probabilistic choice. Thus programs become functions from initial state to *sets* of results, in this case to sets of discrete sub-probability distributions over a (finite) state space.

Adding nondeterminism specialises general powerdomain techniques, but with one novel aspect — namely that we impose (probabilistic) convexity,¹ which we need to allow a refinement relation between demonic nondeterministic choice and probabilistic choice.

For a logic of total correctness over such a model, one might concentrate (as for standard program logic) only on statements that hold with probability

¹This is distinct from the term applied in the construction of the Plotkin powerdomain and is explained below.

1, and for example Rao's logic [20] is based on that idea. However such a logic loses most of the detail of the probabilistic information. To obtain a more discriminating logic Kozen instead used 'expectations' to investigate the final distributions. Briefly his idea is as follows.

Predicates embed into the space of real-valued functions of the state space by identifying *true* and *false* with the reals 1 and 0 respectively; the predicates become 'characteristic functions'. The probability of a program's establishing a postcondition is then the expected value² of the characteristic function over the distribution generated by executing the program: for program *Prog*, initial state *s* and postcondition *post* as a characteristic function we write that *wp.Prog.post.s*.

Once demonic choice is added, expectation transformers are still applicable, yet the interpretation of *wp.Prog.post.s* is then the 'greatest guaranteed' probability that *post* holds after execution from *s*.

We now review both the operational and the expectation-transformer models of probabilistic demonic programs; we state the (partial) Galois connection that exists between them, and the consequences of it. In the expectation logic we use arithmetic operators, very often binary operators lifted pointwise to expectations, such as minimum \sqcap , maximum \sqcup and *p*-averaging:

$$a \oplus_p b := p \times a + (1-p) \times b .$$

We will see that those operators in particular correspond to demonic, angelic and probabilistic choice.

2.2 The operational model

The operational view of probabilistic programs is based on distributions over states.

Definition 2.1 For finite state space *S* the space of *discrete sub-probability distributions* (\bar{S}, \sqsubseteq) over *S* is defined as follows:

$$\bar{S} := \{F: S \rightarrow [0, 1] \mid \sum_{s:S} F.s \leq 1\} ,$$

where $F \sqsubseteq F'$ iff $(\forall s: S \cdot F.s \leq F'.s)$.

Thus the distributions are functions from *S* into the unit interval $[0, 1]$ that sum to no more than 1, and the order \sqsubseteq is just \leq extended pointwise. \square

²The expectation of a $\{0, 1\}$ -valued function *A* is the probability that the event 'A evaluates to 1' occurs.

Note that the distributions do not necessarily sum to 1 exactly; the ‘deficit’ represents the probability of nontermination.³

Now the Smyth powerdomain over \overline{S} is characterised by non-empty up-closed sets of distributions: we say that a subset U of \overline{S} is *up-closed* if whenever F is in U and $F \sqsubseteq F'$ then also F' is in U . Up-closure is one of our conditions on result sets.

Another condition is Cauchy closure:

Definition 2.2 A set $A \subseteq \overline{S}$ is *Cauchy-closed* iff it is a closed subset of \mathbb{R}^S considered as an $|S|$ -dimensional space with the Euclidean topology, where $|S|$ denotes the cardinality of S . \square

Cauchy closure will guarantee our continuity requirement — a recursive program is the limit of the nested result sets that represent its partial unfoldings, which limit must be nonempty. That is guaranteed by the fact⁴ that the intersection of a nested family of nonempty closed sets of a compact topological space is nonempty; but we defer discussion of the compactness of \overline{S} until Sec. 6.

Before we define our operational model we need one more closure condition — the novel *convexity* referred to above — that is not supplied by the standard Smyth construction: it is specific to the space of distributions.

Definition 2.3 A set U in \overline{S} is said to be *convex* if for any real $p \in [0, 1]$ and for any F, F' in U we have also

$$F \text{ }_p\oplus\text{ } F' \in U ,$$

where the scalar operator $\text{ }_p\oplus$ is lifted pointwise to real-valued functions: thus

$$(F \text{ }_p\oplus\text{ } F').s = p \times (F.s) + (1-p) \times (F'.s) .$$

\square

Convex closure is imposed so that a demonic choice made between two programs is ‘refined’ (defined below) by a probabilistic choice, and we justify

³An alternative though less convenient way to model nontermination is to introduce a special bottom state.

⁴We use the finite intersection lemma, a result from topology, which ensures that the intersection of a family of closed sets contained within a compact set is non-empty provided that the intersection of any finite subset of the family is non-empty.

```

var  $s: 1..N$ ;
do
  break  $\frac{1}{s} \oplus s := s - 1$ 
od

```

This deterministic program takes each initial state s in the finite state space $1..N$ to a single distribution which is ‘choose uniformly from $1..s$ ’. (The **do**...**od** construction means ‘loop until **break** is executed’.)

We say ‘deterministic’ — although probabilistic nondeterminism is present — because there is no demonic or angelic nondeterminism in them; more importantly, such programs, like standard deterministic terminating programs, are maximal in the refinement order \sqsubseteq (if miracles [11, 18] are excluded).

Figure 1: Deterministic program over finite state space

that by reasoning that a demonic choice is certainly no better (and in fact probably worse) than one determined by the flip of a coin: in the former case nothing about the method of choosing is known whereas in the latter the distribution is.⁵

Definition 2.4 The space of *probabilistic, demonic programs* $(\mathcal{HS}, \sqsubseteq)$ is given by

$$\mathcal{HS} := S \rightarrow \mathcal{SS} ,$$

where \mathcal{SS} is the set of nonempty, up-closed, convex and Cauchy-closed subsets of \bar{S} , and the refinement order \sqsubseteq is induced pointwise from the Smyth order on \mathcal{SS} . Thus for r, r' in \mathcal{HS} we have

$$r \sqsubseteq r' := (\forall s: S \cdot r.s \supseteq r'.s) .$$

□

That concludes the operational model; Fig. 1 gives an example of a deterministic program.

⁵This is consistent with an assumption that the probability distributions are unaffected by environmental conditions.

2.3 The programming logic

We now turn to the probabilistic logic [17, 7]. Denoting the non-negative reals by \mathbb{R}_{\geq} , we replace predicates (Boolean-valued functions of S) by *bounded expectations* of type $S \rightarrow \mathbb{R}_{\geq}$, denoted $\mathcal{P}S$, where by “bounded” expectation we mean that there is some real number which it nowhere exceeds.⁶ We relate expectations by *probabilistic implication* \Rightarrow defined (with its variants) over $\mathcal{P}S$ by

$$\begin{aligned} \Rightarrow & \text{ ‘everywhere no more than’} \\ \equiv & \text{ ‘everywhere equal to’} \\ \Leftarrow & \text{ ‘everywhere no less than’ .} \end{aligned}$$

We define the space of \Rightarrow -monotonic expectation transformers as follows:

Definition 2.5 The space of *monotonic expectation transformers* $(\mathcal{T}S, \sqsubseteq)$ is given by the set of monotone functions

$$\mathcal{T}S := \mathcal{P}S \rightarrow \mathcal{P}S ,$$

where for t, t' in $\mathcal{T}S$

$$t \sqsubseteq t' \quad \text{iff} \quad (\forall \alpha : \mathcal{P}S \cdot t.\alpha \Rightarrow t'.\alpha) .$$

□

The \Rightarrow symbol reminds us of Boolean implication \Rightarrow since *false* corresponds to the everywhere 0 expectation, and *true* to everywhere 1, and just as *false* \Rightarrow *true* we have $0 \Rightarrow 1$.⁷

The expectation-transformer model generalises the predicate-transformer model for standard (non-probabilistic) demonic programs: programs take ‘post-expectations’ to ‘greatest pre-expectations’. An informal introduction to the logic and its use is given elsewhere [14, 12]. As an example, note that the program of Fig. 1 — call it *ProgA* — satisfies $wp.ProgA.\{1\}.s = 1/s$, that is that the probability of final state $s = 1$ is just the reciprocal of the initial value of s .

⁶For finite state space S the expectations are bounded trivially of course; but in Sec. 6, where we deal with infinite S , boundedness must be imposed.

⁷An immediate benefit is for example the conventional appearance of Def. 2.5 when compared with the definition of refinement between standard transformers t, t'

$$t \sqsubseteq t' \quad \text{iff} \quad (\forall A : \mathbb{P}S \cdot t.A \Rightarrow t'.A) ,$$

where $\mathbb{P}S$ is the set of predicates over S .

2.4 The (partial) Galois connection

We now turn to the partial Galois connection between \mathcal{HS} and \mathcal{TS} : it too generalises a standard construction. ‘Relational programs’ in \mathcal{HS} can be embedded into the transformer space \mathcal{TS} , which is strictly richer; and ‘transformer programs’ can be projected back into the relational space, possibly moving up the \sqsubseteq -order as a result.

We begin with the embedding, a mapping from \mathcal{HS} to \mathcal{TS} .

Definition 2.6 Let h be a program in \mathcal{HS} , so taking initial states in S to sets of final distributions over S . Then the *greatest pre-expectation* at state s of program h , with respect to post-expectation α in \mathcal{PS} , is defined

$$wp.h.\alpha.s \quad := \quad (\sqcap F: h.s \cdot \int_F \alpha),$$

where $\int_F \alpha$ denotes the expected value⁸ of the real-valued function α of S over the distribution F on S , and \sqcap selects the greatest lower bound. We call transformer $wp.h$ the *embedding* of h . \square

Note from the definition that any upper bound for α is an upper bound for $wp.h.\alpha$ also, since $\int_F \alpha$ can be no more than $\sqcup \alpha$; thus boundedness is preserved by $wp.h$ even if S is infinite.

As discussed in the introduction, in the special case when an expectation A is $\{0, 1\}$ -valued, the pre-expectation $wp.r.A.s$ is the greatest guaranteed probability that A holds after execution of r from s — thus with that rationalisation we have replaced ‘holds with certainty’ (as in standard wp) by ‘holds with probability’.

Now we describe the projection, taking expectation transformers in \mathcal{TS} back to relational programs in \mathcal{HS} :

Definition 2.7 For expectation transformer t in \mathcal{TS} the least program $rp.t$ in \mathcal{HS} such that $t \sqsubseteq wp.(rp.t)$ is defined, if it exists, by

$$rp.t.s \quad := \quad \{F: \bar{S} \mid (\forall \alpha: \mathcal{PS} \cdot t.\alpha.s \leq \int_F \alpha)\}.$$

We call $rp.t$ the *projection* of t . \square

⁸In fact $\int_F \alpha$ is just $\sum_{s:S} \alpha.s \times F.s$ because S is finite and F is discrete [3]. Although the expression treats F and α symmetrically in the finite case (a dot product), in the infinite case even the types of F and α are different — and for that reason we use the \int -notation; it is also less cluttered.

The first main result concerning rp and wp is that they establish a partial Galois correspondence between \mathcal{HS} and \mathcal{TS} . We restate it here.

Theorem 2.8 [17, Lemm. 8.2, 8.3] \mathcal{HS} and \mathcal{TS} are put in partial Galois correspondence by rp and wp . Whenever rp is defined, the following inequations hold:

$$\begin{aligned} id &\sqsubseteq wp \circ rp \\ rp \circ wp &= id, \end{aligned}$$

where id is the identity mapping in the respective spaces. \square

The partiality of rp in Def. 2.7 is caused by the possibility that the set

$$\{F: \bar{S} \mid (\forall \alpha: \mathcal{PS} \cdot t.\alpha.s \leq \int_F \alpha)\}$$

might be empty, and we have excluded for this presentation empty results sets from \mathcal{HS} : they correspond to ‘miracles’ [11, 18]. An alternative would be to allow miracles, and the Galois connection would become total; we would then adjoin ∞ to \mathbb{R}_{\geq} , to make it and \mathcal{PS} become \sqcup -complete.⁹

The importance of Thm. 2.8 lies in its corollary, and it forms the second main result. That result establishes ‘healthiness’ conditions in the style of Dijkstra which identify the wp -images of \mathcal{HS} : they identify the subset of ‘implementable’ programs embedded within \mathcal{TS} . The result makes clear when theorems proved in \mathcal{TS} correspond to theorems about programs in \mathcal{HS} , which is important since (as for standard predicate transformers) \mathcal{TS} is a more convenient framework for calculation than \mathcal{HS} is.

The healthiness conditions may be described as three properties of expectation transformers which together characterise the wp -images of programs (Thm. 2.9 below). For a non-negative real a we define scalar multiplication lifted to expectations (for example $a \times \alpha$ in (1)) by pointwise multiplication; similarly we lift the binary operator \ominus on reals defined $a \ominus b := (a - b) \sqcup 0$, where \sqcup is the binary maximum operator; and for a real a we denote by \underline{a} the constant expectation everywhere evaluating to a . The three properties are then

$$\begin{aligned} t.(a \times \alpha) &\equiv a \times t.\alpha && \text{scaling} \\ t.(\alpha + \alpha') &\Leftarrow t.\alpha + t.\alpha' && \text{sub-additivity} \\ t.(\alpha \ominus \underline{a}) &\Leftarrow t.\alpha \ominus \underline{a} && \ominus\text{-subdistribution} . \end{aligned} \tag{1}$$

⁹We do not make those extensions here because ∞ would complicate the arithmetic in our principal use of the Galois connection in this report, which is to establish healthiness conditions for the subspaces of \mathcal{TS} .

We refer collectively to the properties (1) as ‘sublinearity’.¹⁰

The scaling property is particular to an expectation model of probability: if an expectation is scaled by a constant a then the expected value with respect to a distribution is similarly scaled. The two subdistribution properties, on the other hand, together generalise standard conjunctivity [17, following Lem. 7.2].

The importance of sublinearity is that it exactly characterises *wp*-images:

Theorem 2.9 [17, Thm. 8.7] Given any t in \mathcal{TS} , there is an r in \mathcal{HS} with $t = wp.r$ exactly when t is sublinear (1). \square

We denote the *wp*-images of \mathcal{HS} , or equivalently by Thm. 2.9 the subset of sublinear transformers, by $\mathcal{T}_{\sqcap}S$; we shall call them *demonic* expectation transformers,¹¹ and we return to them in Sec. 4.

That concludes the review of our earlier treatment [17] of demonic probabilistic programs over finite state spaces.

3 Deterministic programs

Standard deterministic programs deliver a single output state for any fixed input; probabilistic deterministic programs deliver a single output *distribution* for any fixed input. (Recall Fig. 1.) Thus even though separate runs of a deterministic, probabilistic program may deliver different outputs from the same initial state, over a group of runs a tabulation of results will reveal a probability distribution of those outputs; and because the program is deterministic, a second group of runs from the same initial state will reveal the same distribution.

Since standard deterministic programs are disjunctive as well as conjunctive — and those properties are (\neg) -duals — we conjecture that probabilistic

¹⁰These properties can be more concisely (and correspondingly more cryptically) expressed as follows [17, Def. 7.1]: expectation transformer t is sublinear if and only if for all non-negative scalars a, b, c and expectations α, β we have

$$t.(a \times \alpha + b \times \beta \ominus \underline{c}) \Leftarrow a \times t.\alpha + b \times t.\beta \ominus \underline{c}.$$

Note that $+$ binds more tightly than \ominus .

¹¹Earlier [17] we called these the ‘regular’ transformers, because they correspond to relations.

programs are *super-* as well as sub-additive, since those properties are ($-$)-duals: that is, deterministic programs are simply additive. Our main result in this section is to show that additivity, scaling and \ominus -subdistribution (referred to collectively as ‘linearity’ below) indeed characterise determinism.

We begin by fixing the idea of a deterministic operational program — it is one that is characterised by a single output distribution.

Definition 3.1 A program h in \mathcal{HS} is *deterministic* iff for all states s there is a single distribution F_s so that

$$h.s \quad := \quad \{F: \overline{S} \mid F_s \sqsubseteq F\} .$$

□

Notice how nonterminating programs are also covered by the definition. Deterministic here means “deterministic if terminating” and is thus closer to the notion of pre-deterministic used by some authors [4].

Next we define linearity.

Definition 3.2 An expectation transformer t is said to be *linear* iff it is \ominus -subdistributive, scaling and *additive*, where for that last we mean that for all expectations α, α' in \mathcal{PS} we have

$$t.(\alpha + \alpha').s \quad = \quad t.\alpha.s + t.\alpha'.s .$$

□

The proof that deterministic programs map to linear expectation transformers under wp is straightforward.

Lemma 3.3 If h in \mathcal{HS} is deterministic then $wp.h$ is linear.

Proof: Thm. 2.9 ensures that $wp.h$ is sublinear at least; then its additivity follows directly from Def. 3.1 and the additivity of \int_F (equivalently the additivity of expectation in elementary probability [3]). □

More involved is the converse, that linearity implies determinism.

Lemma 3.4 If a transformer t in \mathcal{TS} is linear, then it is the wp -image of a deterministic relational program.

Proof: For expectation transformer t in \mathcal{TS} we define the deterministic program h in \mathcal{HS} as follows:

$$h.s := \{F: \bar{S} \mid F_s \sqsubseteq F\}$$

where $F_s.s' := t.\{s'\}.s$,

writing $\{s'\}$ for the ‘point’ expectation that is 1 at s' and 0 elsewhere.

Observe that h is well defined: F_s is indeed an element of \bar{S} because

$$\begin{aligned} & \sum_{s':S} F_s.s' \\ = & t.(\sum_{s':S} \{s'\}).s && t \text{ additive; definition } F \\ = & t.\underline{1}.s && \text{definition } \underline{1} \\ \leq & (t.\underline{1}.s \ominus 1) + 1 && \text{arithmetic} \\ \leq & t.\underline{0}.s + 1 && t \text{ is } \ominus\text{-subdistributive} \\ = & 0 \times t.\underline{0}.s + 1 && t \text{ scaling} \\ = & 1 . \end{aligned}$$

We now show that in fact $t = wp.h$. Observe first that any expectation α in \mathcal{PS} is a (finite) sum $\sum_{s:S} (\alpha.s) \times \{s\}$ of characteristic functions. (The sum is finite because S is.) We reason

$$\begin{aligned} & t.\alpha.s \\ = & t.(\sum_{s':S} \alpha.s' \times \{s'\}).s \\ = & \sum_{s':S} \alpha.s' \times t.\{s'\}.s && t \text{ linear} \\ = & \sum_{s':S} \alpha.s' \times F_s.s' && \text{definition of } F_s \\ = & \int_{F_s} \alpha && \text{definition of } \int \\ = & wp.h.\alpha.s . && \int \text{ monotonic; Lem. 3.4 of } h; \text{ Def. 2.6} \end{aligned}$$

□

We call the space of linear — equivalently deterministic — transformers $\mathcal{T}_o S$, and have our theorem:

Theorem 3.5 A transformer is linear iff it is the wp -image of a deterministic relational program.

Proof: Lemm. 3.3, 3.4. □

Fig. 2 illustrates the result.

Let $ProgB$ be the deterministic program

$$\begin{aligned} & \mathbf{var} \ s: \{Heads, Tails\}; \\ & s := Heads \frac{1}{2} \oplus s := Tails . \end{aligned}$$

We have $wp.ProgB.\{Heads\}.s = 1/2$ for all initial values of s , and similarly $wp.ProgB.\{Tails\}.s = 1/2$. Now (punning sets and characteristic functions) we have $\{Heads\} + \{Tails\} = \{Heads, Tails\}$, and

$$wp.ProgB.\{Heads, Tails\}.s = 1 = 1/2 + 1/2 ,$$

thus illustrating additivity.

Figure 2: Deterministic programs are additive

4 Demonic programs

Demonic programs are constructed by closing $\mathcal{T}_\circ S$ under demonic choice \sqcap : we show here that they are exactly the sublinear transformers $\mathcal{T}_\sqcap S$ described in Sec. 2, which were the subject of our earlier work [17]. Fig. 3 shows a ‘typical’ demonic program: first natural number s is decreased demonically from its initial value, but not below 1; then a uniform probabilistic choice is made from the numbers $1..s$. The final distribution over a set of runs is not (wholly) predictable, because of the nondeterminism: in separate sets one could find quite different distributions. But each distribution will have the property that higher frequencies are associated with lower final states; and it can be shown that the program is capable of producing all such ‘anti-monotonic’ distributions.

To establish our characterisation of $\mathcal{T}_\sqcap S$, we note first from simple arithmetic that if t and t' are additive then $t \sqcap t'$ is sub-additive (by subdistribution of greatest lower bounds through addition), and that scaling and \ominus -subdistribution are preserved. Thus the elements of \sqcap -closure of $\mathcal{T}_\circ S$ are sublinear.

For the converse, we note that if t is sublinear then from Thm. 2.9 and Thm. 2.8 we have $wp.(rp.t) = t$,¹² thus take for the deterministic ‘components’ of t the programs in $\mathcal{T}_\circ S$ generated (Def. 3.1) by choosing F_s in all

¹²This is our principal use of the Galois connection: note that sublinearity of t implies it is not miraculous [17], and that in turn ensures definedness of rp .

```

var  $s: 1..N$ ;

do  $s \neq 1 \rightarrow$       /* choose  $s$  demonically from  $1..s$  */
  break  $\sqcap s := s - 1$ 
od;

do      /* choose  $s$  uniformly from  $1..s$  */
  break  $\frac{1}{s} \oplus s := s - 1$ 
od

```

This demonic program takes an initial state s in $1..N$ to any ‘anti-monotonic’ distribution over $1..s$ that assigns higher probabilities to lower outcomes.

Figure 3: Demonic program over finite state space

possible ways from $rp.t.s$ for each s . The result is then a consequence of the correspondence between \cup in \mathcal{HS} and \sqcap in \mathcal{TS} . Thus any sublinear transformer may be written as the (possibly infinite) demonic choice \sqcap between members of $\mathcal{T}_\circ S$, and we have proved this theorem:

Theorem 4.1 The \sqcap -closure of the deterministic transformers $\mathcal{T}_\circ S$ is exactly the sublinear transformers $\mathcal{T}_\sqcap S$. \square

Fig. 4 gives an illustration of sublinearity.

5 Angelic/demonic programs

Our final step is to move from $\mathcal{T}_\sqcap S$ to the more general space formed by allowing angelic choice as well. In the standard case one drops conjunctivity; here we will drop sub-additivity.

Our generalisation of angelic choice in the probabilistic domain is as follows:

Definition 5.1 For expectation transformers t, t' in \mathcal{TS} , expectation α in \mathcal{PS} and state s in S we define

$$(t \sqcup t').\alpha.s := t.\alpha.s \sqcup t'.\alpha.s .$$

Let $ProgC$ be the program

$$\begin{aligned} & \mathbf{var} \ s: \{Heads, Tails\}; \\ & s := Heads \sqcap s := Heads \frac{2}{3} \oplus s := Tails . \end{aligned}$$

This program contains both probabilistic and demonic choice; it chooses between $Heads$ and $Tails$, assigning a probability of *at least* $2/3$ to the former. Now we have $wp.ProgC.\{Heads\}.s = 2/3$ because the nondeterminism is demonic (not angelic) — although $Heads$ can be chosen always, we can be *sure* it will be chosen only two-thirds of the time. Similarly $wp.ProgC.\{Tails\}.s = 0$ because $Tails$ might never be chosen at all. Now

$$wp.ProgC.\{Heads, Tails\}.s = 1 > 2/3 + 0 ,$$

thus illustrating sublinearity and the failure of additivity.

Figure 4: Demonic programs are sublinear but not additive

□

We denote the closure of $\mathcal{T}_\sqcap S$ under \sqcup by $\mathcal{T}_{\sqcup} S$ (using \sqcup for least upper bound in the semantics); its characterisation will turn out to be ‘semi-sublinearity’.

Definition 5.2 Expectation transformer t in $\mathcal{T}S$ is said to be *semi-sublinear* iff it satisfies the properties of scaling and \ominus -subdistribution but not (necessarily) sub-additivity:

$$\begin{aligned} t.(a \times \alpha) & \equiv a \times (t.\alpha) && \text{scaling} \\ t.(\alpha \ominus \underline{a}) & \Leftarrow t.\alpha \ominus \underline{a} && \ominus\text{-subdistribution} . \end{aligned}$$

□

We show first that all members of $\mathcal{T}_{\sqcup} S$ are semi-sublinear.

Lemma 5.3 If t is in $\mathcal{T}_{\sqcup} S$ then it is semi-sublinear.

Proof: We need only show that semi-sublinearity is preserved by (possibly infinite) applications of \sqcup , and that is immediate from arithmetic and Def. 5.2. □

Our major goal however is to show that $\mathcal{T}_{\sqcup} S$ comprises *exactly* the semi-sublinear transformers, and for that we need the converse of Lem. 5.3: that all semi-sublinear transformers in $\mathcal{T}S$ are members of $\mathcal{T}_{\sqcup} S$.

Motivated by Back and von Wright’s construction [1, Lemma 8] for standard programs, given a semi-sublinear t we consider all post-expectations β and express t as the angelic choice over a family t_β of transformers in $\mathcal{T}_\sqcap S$. Each transformer t_β will be the \sqsubseteq -weakest transformer that itself is monotonic and semi-sublinear¹³ and agrees with t at β : because t_β is the weakest such, it will be everywhere no more than t ; and because each t_β agrees with t at β , their supremum will be at least t . The surprise will be that the t_β ’s turn out to be fully (rather than just semi-) sublinear, making them members of $\mathcal{T}_\sqcap S$ as required.

The appropriate definition of t_β is the following:

Definition 5.4 For t a semi-sublinear transformer in $\mathcal{T}S$ and expectation β in $\mathcal{P}S$, let the β -component t_β of t be defined by

$$t_\beta.\alpha.s \quad := \quad (\sqcup c, c': \mathbb{R}_\geq \mid c \times \beta - \underline{c}' \Rightarrow \alpha \cdot c \times (t.\beta.s) - c') ,$$

for α in $\mathcal{P}S$ and s in S . □

To see that t_β is well defined, by which we mean that when applied to α in $\mathcal{P}S$, the result is nowhere infinite nor negative, we appeal first to Lem. 5.6 below to see that $t_\beta.\alpha$ must be everywhere finite; to see that the result is non-negative we need only put c, c' to be 0, 0 in Def. 5.4 and then notice that the definition maximises.

We now prove the three crucial properties of β -components.

The first property is that t_β is no less than t at β .

Lemma 5.5 For all semi-sublinear transformers t and expectations β we have $t.\beta \Rightarrow t_\beta.\beta$.

Proof: Take $c, c' := 1, 0$ in Def. 5.4. □

The second property is that t_β is no more than t in general.

¹³The impact of requiring semi-sublinearity in t_β is that the otherwise obvious choice

$$t_\beta.\alpha \quad := \quad t.\beta \text{ if } (\beta \Rightarrow \alpha) \text{ else } 0 \tag{2}$$

is not allowed — by scaling, for example, we see that $t_\beta.(2\beta)$ must be $2(t.\beta)$ rather than the $t.\beta$ that (2) would give. Those properties of t_β motivate Def. 5.4 that makes it monotonic and semi-sublinear by construction, leaving the job of Lem. 5.7 mainly to establish its sub-additivity.

Lemma 5.6 For all semi-sublinear transformers t and expectations β we have $t_\beta \sqsubseteq t$.

Proof: Take arbitrary expectation α in \mathcal{PS} and state s . For all c, c' in \mathbb{R}_\geq ,

$$\begin{array}{ll}
& c \times \beta - \underline{c}' \Rightarrow \alpha \\
\text{implies} & c \times \beta \ominus \underline{c}' \Rightarrow \alpha & \underline{0} \Rightarrow \alpha \\
\text{implies} & t.(c \times \beta \ominus \underline{c}').s \leq t.\alpha.s & t \text{ monotonic} \\
\text{implies} & c \times (t.\beta.s) \ominus \underline{c}' \leq t.\alpha.s & t \text{ semi-sublinear} \\
\text{implies} & c \times (t.\beta.s) - \underline{c}' \leq t.\alpha.s & \text{arithmetic}
\end{array}$$

Since the inequality holds for arbitrary c, c' it holds for the supremum taken over them in Def. 5.4. \square

Finally we show sublinearity of the β -components.

Lemma 5.7 For semi-sublinear t in \mathcal{TS} and β in \mathcal{PS} the component t_β is sublinear.

Proof: The scaling of t_β is obvious from its definition. The \ominus -subdistributivity is straightforward also:

$$\begin{array}{ll}
& t_\beta.(\alpha \ominus \underline{a}).s \\
= & (\sqcup c, c': \mathbb{R}_\geq \mid c \times \beta - \underline{c}' \Rightarrow \alpha \ominus \underline{a} \cdot c \times (t.\beta.s) - \underline{c}') & \text{Def. 5.4} \\
\geq & & \text{take } c' := c'' + a \\
& (\sqcup c, c'': \mathbb{R}_\geq \mid c \times \beta - (\underline{c}'' + a) \Rightarrow \alpha \ominus \underline{a} \cdot c \times (t.\beta.s) - (c'' + a)) \\
\geq & (\sqcup c, c'': \mathbb{R}_\geq \mid c \times \beta - \underline{c}'' \Rightarrow \alpha \cdot (c \times (t.\beta.s) - c'') - a) & \text{arithmetic} \\
\geq & (\sqcup c, c'': \mathbb{R}_\geq \mid c \times \beta - \underline{c}'' \Rightarrow \alpha \cdot (c \times (t.\beta.s) - c'')) - a & \text{arithmetic} \\
= & t_\beta.\alpha.s - a, & \text{Def. 5.4}
\end{array}$$

and hence, since $t_\beta.(\alpha \ominus \underline{a})$ is everywhere non-negative, we can deduce that indeed

$$t_\beta.(\alpha \ominus \underline{a}).s \geq (t_\beta.\alpha \ominus \underline{a}).s .$$

The main point however is sub-additivity, for which we argue

$$\begin{array}{ll}
& t_\beta.(\alpha_1 + \alpha_2).s \\
= & (\sqcup c, c': \mathbb{R}_\geq \mid c \times \beta - \underline{c}' \Rightarrow \alpha_1 + \alpha_2 \cdot c \times (t.\beta.s) - \underline{c}') & \text{Def. 5.4} \\
= & (\sqcup c_1, c_2, c'_1, c'_2: \mathbb{R}_\geq & \text{take } c, c' := c_1 + c_2, c'_1 + c'_2 \\
& \mid (c_1 + c_2)\beta - (\underline{c}'_1 + \underline{c}'_2) \Rightarrow \alpha_1 + \alpha_2 \\
& \cdot (c_1 + c_2)(t.\beta.s) - (c'_1 + c'_2))
\end{array}$$

$$\begin{aligned}
&\geq && (\sqcup c_1, c_2, c'_1, c'_2: \mathbb{R}_{\geq} && \text{arithmetic} \\
&&& | c_1 \times \beta - \underline{c'_1} \Rightarrow \alpha_1 \wedge c_2 \beta - \underline{c'_2} \Rightarrow \alpha_2 \\
&&& \cdot (c_1 \times (t.\beta.s) - c'_1) + (c_2 \times (t.\beta.s) - c'_2)) \\
&= && && c_1, c'_1 \text{ and } c_2, c'_2 \text{ vary independently} \\
&&& (\sqcup c_1, c'_1: \mathbb{R}_{\geq} | c_1 \times \beta - \underline{c'_1} \Rightarrow \alpha_1 \cdot c_1 \times (t.\beta.s) - c'_1) \\
&+ && (\sqcup c_2, c'_2: \mathbb{R}_{\geq} | c_2 \times \beta - \underline{c'_2} \Rightarrow \alpha_2 \cdot c_2 \times (t.\beta.s) - c'_2) \\
&= && t_\beta.\alpha_1.s + t_\beta.\alpha_2.s . && \text{Def. 5.4}
\end{aligned}$$

□

That gives our main theorem for this section:

Theorem 5.8 The \sqcup -closure of the demonic transformers $\mathcal{T}_{\sqcup} S$ is exactly the semi-sublinear transformers $\mathcal{T}_{\square} S$.

Proof: For semi-sublinear expectation transformer t and expectation α in $\mathcal{P}S$ we have

$$\begin{aligned}
&&& t.\alpha \\
\Rightarrow &&& t_\alpha.\alpha && \text{Lem. 5.5} \\
\Rightarrow &&& (\sqcup \beta: \mathcal{P}S \cdot t_\beta).\alpha \\
\Rightarrow &&& t.\alpha . && \text{Lem. 5.6}
\end{aligned}$$

Thus we have $t = (\sqcup \beta: \mathcal{P}S \cdot t_\beta)$, and conclude from Lem. 5.7 that indeed t is the supremum of sublinear transformers.

The converse implication is given by Lem. 5.3. □

Figures 5 and 6 illustrate the behaviour of transformers in $\mathcal{T}_{\square} S$.

6 Infinite state spaces

Our earlier results about $\mathcal{H}S$ and $\mathcal{T}S$ [17] were summarised in Sec. 2, and assumed that the state space is finite; the subsequent analysis of Sections 3–5 also assumed finiteness. We now show how to extend all those results to infinite state spaces; the principal tool will be continuity.

In summary there will be two changes: the first is that we must explicitly restrict expectations to $\mathcal{P}_B S$, the subset of $\mathcal{P}S$ consisting of the *bounded*

Let $ProgD$ be the program

$$\begin{aligned} & \mathbf{var} \ s: \{Heads, Tails\}; \\ & s := Heads \frac{1}{3} \oplus s := Tails \quad \sqcap \quad s := Tails, \end{aligned}$$

and take $ProgC$ as in Fig. 4. Programs $ProgC$ and $ProgD$ are in $\mathcal{T}_{\sqcap} S$; let $ProgE$ in $\mathcal{T}_{\sqcap} S$ be the angelic choice $ProgC \sqcup ProgD$ between them.

Program $ProgE$ (confusingly) chooses *angelically* whether to use a *Heads*-biased or *Tails*-biased coin, and then applies the chosen bias demonically.

Now we have $wp.ProgE.\{Heads\}.s = 2/3 \sqcup 0 = 2/3$ because the nondeterminism is angelic between the biases, and once the *Heads* bias is chosen, the worst it can do is to select *Heads* only two-thirds of the time. Similarly $wp.ProgE.\{Tails\}.s = 2/3$. Yet

$$wp.ProgE.\{Heads, Tails\}.s = 1 < 2/3 + 2/3,$$

thus illustrating the failure of sublinearity.

Figure 5: Angelic programs are not sublinear

Let $ProgF$ be the program

$$\mathbf{skip} \frac{2}{3} \oplus \mathbf{abort}.$$

This deterministic program takes any initial state s to the final distribution assigning $2/3$ to s and 0 to all other states. The ‘deficit’ $1 - 2/3$ is the probability $1/3$ of nontermination caused by the primitive non-terminating program **abort**. Now have $wp.ProgF.1.s = 2/3$ for all s , and indeed $wp.ProgF.r.s = 2r/3$ for any real number r , illustrating scaling. An illustration of \ominus -subdistribution is that

$$wp.ProgF.(1 \ominus 1/2).s = 2/3 \times (1 \ominus 1/2) > 1/6 = 2/3 \ominus 1/2.$$

Figure 6: Angelic programs are semi-sublinear

functions from S into \mathbb{R}_{\geq} ; and the second is that we must address the issue of continuity explicitly.¹⁴ The restriction to bounded expectations α is necessary to guarantee that $\int_F \alpha$ is well defined (is finite in value)¹⁵. And continuity of $wp.Prog$ for $Prog$ in \mathcal{HS} is no longer to be taken for granted: its derivation from sublinearity [17, Lem. 7.6] depended on finiteness of S . Thus when S is infinite we must prove continuity directly from Def. 2.4 (the operational model) and Def. 2.6 (wp); and that forms the main technical contribution of this section.

Once continuity is dealt with, most of our earlier results [17] can be re-established by appeal to their proofs for the finite case.

We begin by introducing some technical details of the probability distributions over the now-infinite state space.

6.1 Topological preliminaries; compactness of \overline{S}

Our set of distributions \overline{S} is a subset of the function space $S \rightarrow \mathbb{R}_{\geq}$, which in turn can be regarded equivalently as Euclidean space \mathbb{R}_{\geq}^S but now possibly of infinite dimension.¹⁶ The principal notions we use there are topological closure and compactness, the latter being essential for establishing continuity of wp -images.

The *Euclidean topology* for the real line \mathbb{R}_{\geq} corresponds to the usual Euclidean metric, and the topology \mathcal{E}_S we use for \mathbb{R}_{\geq}^S is the product of the Euclidean topologies over \mathbb{R}_{\geq} for each dimension in S . A basis for \mathcal{E}_S is then the collection of sets of the form

$$N \times \mathbb{R}_{\geq}^{S-T} \tag{3}$$

(ignoring order in the product), for all finite subsets T of S and open subsets N of \mathbb{R}_{\geq}^T . For dimension s in S we say that s is in the *support* of an open

¹⁴Finiteness of S gives boundedness and continuity automatically; so these issues do not limit our earlier results in any way.

¹⁵In the infinite case we define $\int_F \alpha$ by

$$(\sqcup T: \mathbb{F}S \cdot \sum_{s:T} F.s \times \alpha.s) .$$

¹⁶One generalisation we do not (need to) make is from discrete distributions over finite S to measures over an infinite S . That is because our programs' only access to probability is via a binary operator $_p\oplus$, and they therefore can generate only discrete result distributions, even when S is infinite.

set in \mathcal{E}_S just when the projection of that set onto s is not all of \mathbb{R}_{\geq} . Similar notions of finiteness apply to expectations:

Definition 6.1 An expectation α in $\mathcal{P}S$ is *finitary* if there is a finite subset T of S such that $\alpha.s \neq 0$ only for states s in T . (Note that such expectations are bounded by construction.)

The *support* of an expectation α is the set of states s such that $\alpha.s \neq 0$; thus an expectation is finitary iff it has finite support.

We write $\mathbb{F}S$ for the set of finite subsets of S , and $\mathcal{P}_f S$ for the set of finitary expectations over S .

For subset T of S , we define the *restriction* of α to T as

$$(\alpha \downarrow T).s \quad := \quad \begin{array}{ll} \alpha.s & \text{if } s \in T \\ 0 & \text{otherwise.} \end{array}$$

Thus $\alpha \downarrow T$ is finitary if T is finite. □

Our technique for extending our results will be to show that they continue to hold in infinite S provided we restrict our attention to finitary expectations; then continuity of the transformers involved will carry the equalities through to all expectations, since any expectation can be written as a directed \sqcup -limit of its finitary restrictions.

First we must investigate some properties of finitary expectations themselves; we begin by showing that finitary expectations can be used to define ‘closed half-spaces’ in \mathbb{R}_{\geq}^S .

Lemma 6.2 For (finitary) α in $\mathcal{P}_f S$ and r in \mathbb{R}_{\geq} , sets of the forms

$$\{F: \mathbb{R}_{\geq}^S \mid \int_F \alpha \leq r\} \quad \text{and} \quad \{F: \mathbb{R}_{\geq}^S \mid \int_F \alpha \geq r\}$$

are closed in the topology \mathcal{E}_S over \mathbb{R}_{\geq}^S .

Proof: Let the support of α be \bar{T} , finite because α is finitary. Then in either case above the projection of the set is all of \mathbb{R}_{\geq} for each dimension outside of T ; and its projection into \mathbb{R}_{\geq}^T is the complement of an open set there. □

For infinitary expectations, however, we have closure in one direction only.

Lemma 6.3 For any expectation α in \mathcal{P}_{BS} and r in \mathbb{R}_{\geq} , the set of distributions

$$\{F: \mathbb{R}_{\geq}^S \mid \int_F \alpha \leq r\}$$

is closed in \mathbb{R}_{\geq}^S .

Proof: We have

$$\begin{aligned} & \{F: \mathbb{R}_{\geq}^S \mid \int_F \alpha \leq r\} \\ = & \{F: \mathbb{R}_{\geq}^S \mid \int_F (\sqcup T: \mathbb{F}S \cdot \alpha \downarrow T) \leq r\} \\ = & \text{bounded monotone convergence for } \int_F \text{ [5]} \\ & \{F: \mathbb{R}_{\geq}^S \mid (\sqcup T: \mathbb{F}S \cdot \int_F \alpha \downarrow T) \leq r\} \\ = & \{F: \mathbb{R}_{\geq}^S \mid (\forall T: \mathbb{F}S \cdot \int_F \alpha \downarrow T \leq r)\} \\ = & (\cap T: \mathbb{F}S \cdot \{F: \mathbb{R}_{\geq}^S \mid \int_F \alpha \downarrow T \leq r\}) , \end{aligned}$$

which is an intersection of a (T -indexed) collection of closed sets (Lem. 6.2), since $\alpha \downarrow T$ is finitary for all T in $\mathbb{F}S$. \square

That Lem. 6.3 works in only one direction is because our expectations (finitary or not) take only non-negative values. To see that its dual does not hold in general, consider the half-space

$$\{F: \mathbb{R}_{\geq}^S \mid \int_F \underline{1} \geq 1\} . \quad (4)$$

The origin (F everywhere 0) does not lie within it, yet every open set in the basis (3) containing the origin also intersects the half-space — for example, take F in \mathbb{R}_{\geq}^S to be 1 at some point in $S-T$ and 0 elsewhere. Thus the origin is a limit point of (4) but is not in it.

With the above we now have the compactness property of our space of distributions referred to earlier.

Lemma 6.4 The space \bar{S} of distributions over S is a compact subset of $(\mathbb{R}_{\geq}^S, \mathcal{E}_S)$.

Proof: The set \bar{S} may be written

$$[0, 1]^S \cap \{F: \mathbb{R}_{\geq}^S \mid \int_F \underline{1} \leq 1\} .$$

But $[0, 1]^S$ is compact by Tychonoff's Theorem¹⁷, and $\{F: \mathbb{R}_{\geq}^S \mid \int_F \underline{1} \leq 1\}$ is closed by Lem. 6.3. Thus \bar{S} is the intersection of a compact set and a closed set, and is therefore compact. (It is also closed.) \square

¹⁷ *Tychonoff's Theorem* states that a product of compact sets is compact in the product topology.

6.2 The Galois functions; continuity

We now return to the connection (Thm. 2.8) between the relational space $\mathcal{H}S$ and the expectation-transformer space $\mathcal{T}S$: they are connected by the functions

$$\begin{aligned} rp: \mathcal{T}S &\rightarrow \mathcal{H}S && \text{(Def. 2.7)} \\ \text{and } wp: \mathcal{H}S &\rightarrow \mathcal{T}S && \text{(Def. 2.6) .} \end{aligned}$$

For well-definedness of rp in the finite case we showed [17] that given any t in $\mathcal{T}S$ (for which rp is defined) and s in S the set $rp.t.s$ was up-closed, convex and Cauchy-closed: and they are just the conditions placed on result sets in $\mathcal{H}S$ by Def. 2.4.

In the infinite case rp is not so well behaved: up-closure and convexity follow as before; but Cauchy-closure now requires an explicit appeal to the continuity of t . Here we are concerned with finite results, thus we use the weaker notion of ‘bounded continuity’; it is also exactly the condition we need to establish our results.

Definition 6.5 A transformer t in $\mathcal{T}S$ is *boundedly continuous* if for any directed subset \mathcal{A} of expectations bounded above by some constant,

$$t.(\sqcup\mathcal{A}) = (\sqcup\alpha: \mathcal{A} \cdot t.\alpha) .$$

□

The following technical lemma shows that when t is boundedly continuous we may restrict ourselves to finitary expectations in the use of Def. 2.7:

Lemma 6.6 For any boundedly-continuous t in $\mathcal{T}S$ and s in S we have

$$rp.t.s = \{F: \bar{S} \cdot (\forall\alpha: \mathcal{P}_f S \cdot t.\alpha.s \leq \int_F \alpha)\} .$$

The only change from Def. 2.7 is the type of the universally quantified α — it is now drawn from $\mathcal{P}_f S$ rather than $\mathcal{P}S$.

Proof: For arbitrary α in $\mathcal{P}_B S$,

$$\begin{aligned} & (\forall T: \mathbb{F}S \cdot t.(\alpha \downarrow T).s \leq \int_F \alpha \downarrow T) \\ \text{implies } & (\sqcup T: \mathbb{F}S \cdot t.(\alpha \downarrow T).s) \leq (\sqcup T: \mathbb{F}S \cdot \int_F \alpha \downarrow T) \end{aligned}$$

iff t boundedly continuous; bounded monotone convergence

$$t.(\sqcup T: \mathbb{F}S \cdot \alpha \downarrow T).s \leq \int_F (\sqcup T: \mathbb{F}S \cdot \alpha \downarrow T)$$

$$\text{iff } t.\alpha.s \leq \int_F \alpha .$$

The result then follows directly from Def. 2.7, since any $\alpha \downarrow T$ is finitary. \square

With Lem. 6.6 we can establish Cauchy-closure of $rp.t$ for boundedly-continuous t .

Lemma 6.7 For any boundedly-continuous t in $\mathcal{T}S$ and state s in S , the set of distributions $rp.t.s$ is Cauchy-closed in \overline{S} .

Proof: We reason

$$\begin{aligned} & rp.t.s \\ = & \{F: \overline{S} \cdot (\forall \alpha: \mathcal{P}_f S \cdot t.\alpha.s \leq \int_F \alpha)\} && \text{Lem. 6.6} \\ = & (\cap \alpha: \mathcal{P}_f S \cdot \{F: \overline{S} \mid t.\alpha.s \leq \int_F \alpha\}) , \end{aligned}$$

which by Lem. 6.2 is an intersection of Cauchy-closed sets. \square

Having shown boundedly-continuous t to yield Cauchy-closed $rp.t$, we turn to the converse: that Cauchy-closed h yields boundedly-continuous $wp.h$. For infinite S it must be done directly, rather than from sublinearity.

Lemma 6.8 For any relational program h in $\mathcal{H}S$, the corresponding expectation transformer $wp.h$ is boundedly continuous.

Proof: Let \mathcal{A} be a \Rightarrow -directed and bounded subset of expectations in $\mathcal{P}_B S$; we show that for any $c > 0$ and state s in S

$$wp.h.(\sqcup \mathcal{A}).s \leq (\sqcup \alpha: \mathcal{A} \cdot wp.h.\alpha.s) + c ,$$

which is sufficient for continuity since c may be arbitrarily small. (The other direction is given by monotonicity.)

Define $r := (\sqcup \alpha: \mathcal{A} \cdot wp.h.\alpha.s)$; then we have

$$\begin{aligned} & (\forall \alpha: \mathcal{A} \cdot wp.h.\alpha.s \leq r) && \text{definition of } r \\ \text{iff} & (\forall \alpha: \mathcal{A} \cdot (\cap F: h.s \cdot \int_F \alpha \leq r)) && \text{Def. 2.6} \\ \text{implies} & (\forall \alpha: \mathcal{A} \cdot (\exists F: h.s \cdot \int_F \alpha \leq r + c)) && c > 0 \\ \text{iff} & (\forall \alpha: \mathcal{A} \cdot \{F: h.s \cdot \int_F \alpha \leq r + c\} \neq \emptyset) \end{aligned}$$

implies Lem. 6.3; $h.s$ closed in \overline{S} ; \mathcal{A} directed; \overline{S} compact — see below

```

var  $s: \mathbb{N}$ ;
 $s := 0$ ;
do
  break  $\frac{1}{2} \oplus s := s + 1$ 
od

```

This deterministic program takes all initial states to the same final distribution, namely one which assigns probability $1/2^{s+1}$ to state s .

Although the program is ‘infinitely branching’ — every natural number is reachable with non-zero probability — it is still \sqcup -continuous since the branching is probabilistic, not demonic or angelic. (Recall that in contrast infinitely \sqcap -branching programs are not generally \sqcup -continuous.)

Figure 7: Continuous program over infinite state space

$$\begin{array}{ll}
& (\cap \alpha: \mathcal{A} \cdot \{F: h.s \cdot \int_F \alpha \leq r + c\}) \neq \emptyset \\
\text{iff} & (\exists F: h.s \cdot (\forall \alpha: \mathcal{A} \cdot \int_F \alpha \leq r + c)) \\
\text{iff} & (\exists F: h.s \cdot \int_F (\sqcup \mathcal{A}) \leq r + c) \quad \text{bounded monotone convergence} \\
\text{implies} & (\sqcap F: h.s \cdot \int_F (\sqcup \mathcal{A})) \leq r + c \\
\text{iff} & wp.h.(\sqcup \mathcal{A}).s \leq r + c. \quad \text{Def. 2.6}
\end{array}$$

For the deferred justification note that Lem. 6.3 and Cauchy-closure of $h.s$ imply Cauchy-closure of each set $\{F: h.s \cdot \int_F \alpha \leq r + c\}$, and \mathcal{A} ’s being directed ensures that they have the finite intersection property. \square

Fig. 7 illustrates continuity for a transformer over an infinite state space.

6.3 The Galois connection

Now we re-establish the partial Galois connection, and with it our characterisation of demonic programs, for infinite state spaces.

Let $\mathcal{T}_C S$ be the boundedly continuous expectation transformers in $\mathcal{T} S$. The fact that rp and wp form a partial Galois connection between $\mathcal{H} S$ and $\mathcal{T}_C S$ is not difficult to show: the inequalities

$$rp \circ wp \sqsubseteq id \quad \text{and} \quad id \sqsubseteq wp \circ rp$$

do not require finiteness of S . However our original proofs of the stronger results, namely the equality $rp \circ wp = id$ and that sublinearity characterises the wp images in \mathcal{TS} (reported above as Thm. 2.9 [17, Lemm. 8.2, 8.6]), did use finiteness.

Our new proof for the first is as follows.

Lemma 6.9 For any h in \mathcal{HS} we have

$$rp.(wp.h) = h .$$

Proof: The proof for the finite case [17, Lem. 8.2] appealed to the *separating hyperplane lemma*¹⁸ [17, Lem. A.1]. Using Lem. A.1 (of this paper) instead allows the proof to go through for the infinite case. \square

To recover the second result we must strengthen its assumptions to include bounded continuity; then we have

Lemma 6.10 If (boundedly-continuous) t in $\mathcal{TC}S$ is sublinear, then

$$wp.(rp.t) = t .$$

Proof: We are able to replay our earlier proof [17, Lem. 8.6] that established $wp.(rp.t).\alpha \equiv t.\alpha$, but do so only for expectations α with finite support. (Full details are given elsewhere [10].)

Bounded continuity of t then allows the result to be extended to all α : on the left we have that $rp.t$ is Cauchy-closed because t is boundedly continuous, and that $wp.(rp.t)$ is boundedly continuous because $rp.t$ is Cauchy-closed; on the right t itself is boundedly continuous. We therefore take limits on both sides over finitary restrictions $\alpha \downarrow T$ of α . \square

6.4 Healthiness conditions

The results above show that in our characterisation of wp -images of programs in \mathcal{HS} , the healthiness conditions Fig. 8 are applicable to the infinite case as well. The first three conditions in Fig. 8 are consequences of sublinearity as shown by our earlier work [17] and, in the special case where the scalars are $\{0, 1\}$ -valued, all are generalisations of properties of predicate transformers.

¹⁸The separating hyperplane lemma is a standard result from linear programming and (specialised to this context) says that given any convex set \mathcal{C} of distributions over a finite state space, and any distribution $F \notin \mathcal{C}$, there is a hyperplane separating F from \mathcal{C} .

<i>feasibility</i>	$t.\alpha \Rightarrow \sqcup \alpha$	for $\alpha: \mathcal{P}_B S$
<i>monotonicity</i>	$\alpha_1 \Leftarrow \alpha_2$ implies $t.\alpha_1 \Leftarrow t.\alpha_2$	for $\alpha_1, \alpha_2: \mathcal{P}_B S$
<i>scaling</i>	$t.(c\alpha) \equiv c(t.\alpha)$	for $\alpha: \mathcal{P}_B S$ and $c: \mathbb{R}_{\geq}$
<i>sublinearity</i>	$c_1(t.\alpha_1) + c_2(t.\alpha_2) \ominus \underline{c_0}$ $\Rightarrow t.(c_1\alpha_1 + c_2\alpha_2 \ominus \underline{c_0})$	for $\alpha_1, \alpha_2: \mathcal{P}_B S$ and $c_0, c_1, c_2: \mathbb{R}_{\geq}$
<i>bounded continuity</i>	$t.(\sqcup \mathcal{A}) \equiv (\sqcup \alpha: \mathcal{A} \cdot t.\alpha)$	for a bounded, \Rightarrow -directed subset \mathcal{A} of $\mathcal{P}_B S$

Figure 8: The properties characterising $\mathcal{T}_{\sqcup} S$ in the infinite case (carried over intact from the finite [17]).

Sublinearity is the appropriate generalisation of conjunctivity [17, following Lem. 7.2], and feasibility, for example, generalises strictness (another of Dijkstra’s original axioms). To see that, we reason first

$$t.\underline{0} \Rightarrow \sqcup \underline{0} \equiv 0 ,$$

and then, noting the inequality $t.\underline{0} \Leftarrow \underline{0}$, we can deduce that in fact $t.\underline{0} \equiv \underline{0}$.

The results of Sections 3–5 also remain valid within the space of boundedly-continuous expectation transformers. All our constructions preserve continuity — however Def. 5.4 must be applied only to finitary expectations, and then Thm. 5.8 establishes any boundedly-continuous and semi-sublinear transformer as a supremum of sublinear ones. As summarised in this section, the details of our proofs where we must appeal to theorems from the theory of convex sets are reduced to reasoning over a finite projection of the state space, where those theorems are still valid.

7 Conclusion

This paper has contributed to the theory of probabilistic programs in two ways: we have explored the structure of the expectation-transformer space (Fig. 9); and we have extended the earlier theory [17], together with the results of the exploration here, to infinite state spaces.

Even with infinite state spaces, we need only use discrete probability distributions. That is because under the assumption [17] that any implementation of a probabilistic specification can only use the discrete choice ${}_p\oplus$ it

On the standard side, the angelic/demonic programs exhaust the space of monotonic transformers. There are however monotonic probabilistic transformers lying outside of $\mathcal{T}_{\square} S$.

Figure 9: Structure of transformer spaces

follows that all implementations will result only in discrete distributions. An interesting topic would be to reinstate continuous distributions [5] into our programs and then to investigate their effects on the healthiness conditions.

In the standard case, the Back and von Wright’s decomposition [1] of the transformer space into deterministic, demonic and angelic components explains the complete behaviour of all monotonic predicate transformers; thus in that theory monotonicity remains as the only surviving restriction — and a necessary one — since in the activity of program specification and development monotonicity is paramount.

For expectation transformers the situation is very different — our Thm. 5.8 implies that there is still part of the space as yet uninvestigated, namely the monotonic transformers that fail to be semi-sublinear. One such expectation transformer is the *weakest liberal* expectation of a program: applied to a standard post-condition it returns the probability that either the post-condition is established, or that the program never terminates — though monotonic, it is not semi-sublinear. (The definitions [9] and applications [12] of weakest-liberal expectations are given elsewhere.) However the general importance of this part of the space in a programming context is far from clear and it remains a topic for further investigation. Other examples of ‘exotic’ expectation transformers are the temporal operators [16, 15] that generalise Kozen’s predicate transformer definitions [8].

The practical import of our results is to give laws in the (arithmetic) programming logic that are satisfied by programs in the various levels of the hierarchy. Such laws are the basis for derived program-development rules such as the use of variants and invariants for loops [12] and, more generally, for programming algebras and even probabilistic temporal logics [16, 15].

A More linear programming lemmas

Lemma A.1 Let \mathcal{C} be a convex subset of \mathbb{R}^S that is compact (hence closed) in the product \mathcal{E}_S of the Euclidean topologies over its constituent projections \mathbb{R} . If some p does not lie in \mathcal{C} , then there is a separating hyperplane with p on one side of it and all of \mathcal{C} on the other.

Proof: If $p \notin \mathcal{C}$, then because \mathcal{C} is closed there is some neighbourhood N in the basis of \mathcal{E}_S with $p \in N$ and $N \cap \mathcal{C} = \emptyset$.

Let T in $\mathbb{F}S$ be the support of N . Writing $(\downarrow T)$ for projection onto T , we then have $p \downarrow T \in N \downarrow T$ and $N \downarrow T \cap \mathcal{C} \downarrow T = \emptyset$, because for the latter

$N\downarrow(S - T) = \mathbb{R}^{S-T}$, and thus $p\downarrow T \notin \mathcal{C}\downarrow T$. Note that $\mathcal{C}\downarrow T$ is compact because \mathcal{C} is, hence closed; and it is convex also because \mathcal{C} is.

Applying the standard separating hyperplane lemma for the finite dimensional space [17, Lem. A1] in the case of $p\downarrow T$ and $\mathcal{C}\downarrow T$, within the finite-dimensional \mathbb{R}^T , gives us a separating hyperplane there; its extension parallel to the remaining axes $S - T$ is then the hyperplane we seek in \mathbb{R}^S . \square

References

- [1] R.-J. R. Back and J. von Wright. Duality in specification languages: a lattice theoretical approach. *Acta Informatica*, 27:583–625, 1990.
- [2] E.W. Dijkstra. *A Discipline of Programming*. Prentice Hall International, Englewood Cliffs, N.J., 1976.
- [3] G. Grimmett and D. Welsh. *Probability: an Introduction*. Oxford Science Publications, 1986.
- [4] Wim H Hesselink. *Programs, Recursion and Unbounded Choice*. Number 27 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, Cambridge, U.K., 1992.
- [5] C. Jones. Probabilistic nondeterminism. Monograph ECS-LFCS-90-105, Edinburgh Univ. Edinburgh, U.K., 1990. (PhD Thesis).
- [6] D. Kozen. Semantics of probabilistic programs. *Journal of Computer and System Sciences*, 22:328–350, 1981.
- [7] D. Kozen. A probabilistic PDL. In *Proceedings of the 15th ACM Symposium on Theory of Computing*, New York, 1983. ACM.
- [8] D. Kozen. Results on the propositional μ -calculus. *Theoretical Computer Science*, 27:333–354, 1983.
- [9] A. K. McIver and C.C. Morgan. Partial correctness for probabilistic programs. Submitted to *Theoretical Computing Science*; available at [19].
- [10] Annabelle McIver and Carroll Morgan. Probabilistic predicate transformers: part 2. Technical Report PRG-TR-5-96, Programming Research Group, March 1996. Available at [19].

- [11] C. C. Morgan. The specification statement. *ACM Transactions on Programming Languages and Systems*, 10(3), July 1988. Reprinted in [13].
- [12] C. C. Morgan. Proof rules for probabilistic loops. In He Jifeng, John Cooke, and Peter Wallis, editors, *Proceedings of the BCS-FACS 7th Refinement Workshop*, Workshops in Computing. Springer Verlag, July 1996.
- [13] C. C. Morgan and T. N. Vickers, editors. *On the Refinement Calculus*. FACIT Series in Computer Science. Springer-Verlag, Berlin, 1994.
- [14] Carroll Morgan. *pGCL: Formal reasoning for random algorithms*. To appear in *South African Computer Journal* in 1999; available at [19].
- [15] Carroll Morgan and Annabelle McIver. An expectation-based model for probabilistic temporal logic. Technical Report PRG-TR-20-97, Programming Research Group, 1997. Available at [19].
- [16] Carroll Morgan and Annabelle McIver. A probabilistic temporal calculus based on expectations. In Lindsay Groves and Steve Reeves, editors, *Proc. Formal Methods Pacific '97*. Springer Verlag Singapore, July 1997. Available at [19].
- [17] Carroll Morgan, Annabelle McIver, and Karen Seidel. Probabilistic predicate transformers. *ACM Transactions on Programming Languages and Systems*, 18(3):325–353, May 1996.
- [18] G. Nelson. A generalization of Dijkstra’s calculus. *ACM Transactions on Programming Languages and Systems*, 11(4):517–561, October 1989.
- [19] PSG. Probabilistic Systems Group: Collected reports.
<http://www.comlab.ox.ac.uk/oucl/groups/probs/bibliography.html>.
- [20] J. R. Rao. Reasoning about probabilistic parallel programs. *ACM Transactions on Programming Languages and Systems*, 16(3), May 1994.