

Compositional Game Theory with Mixed Strategies: Probabilistic Open Games Using a Distributive Law

Neil Ghani
University of Strathclyde
neil.ghani@strath.ac.uk

Clemens Kupke
University of Strathclyde
clemens.kupke@strath.ac.uk

Alasdair Lambert
University of Strathclyde
alasdair.lambert@strath.ac.uk

Fredrik Nordvall Forsberg
University of Strathclyde
fredrik.nordvall-forsberg@strath.ac.uk

We extend the open games framework for compositional game theory to encompass also mixed strategies, making essential use of the discrete probability distribution monad. We show that the resulting games form a symmetric monoidal category, which can be used to compose probabilistic games in parallel and sequentially. We also consider morphisms between games, and show that intuitive constructions give rise to functors and adjunctions between pure and probabilistic open games.

1 Introduction

The research project of open games aims to re-develop the foundations of economic game theory using compositionality and category theory [6], building on e.g. the work of Escardó and Oliva [1]. A compositional framework was proposed by Ghani et al. [3], which included operators from which to build games from smaller component games, and solution concepts such as pure Nash equilibria. However many games that can be found even early on in an undergraduate textbook on game theory (such as e.g. Leyton-Brown and Shoham [10]) fail to contain any equilibria, unless probabilistic (so-called *mixed*) strategies are allowed. In contrast, already Nash [11] proves that mixed strategy Nash equilibria always exist for games with a finite number of players and strategies.

In this work, we extend the framework of open games also to mixed strategies. We use the discrete probability distribution monad on Set (a baby version of the Giry monad [5]) to incorporate probability distributions. However simply moving to the Kleisli category for this monad is not sufficient for our purposes, as that would fail to capture mixed strategies without also demanding e.g. probabilistic play functions. Instead, we make sure to enrich the framework of open games with measured use of the distribution monad in the appropriate places. In particular, we construct a “relational Kleisli lifting”, a variant of the relational lifting for set functors (cf. e.g. Kupke, Kurz and Venema [9]), that turns predicates with non-probabilistic parameters into predicates with mixed parameters in a non-trivial way.

2 Compositional Game Theory with Pure Strategies

We briefly recall the definition of non-probabilistic, “pure” open games as introduced by Hedges [6] for modelling economic game theory with deterministic agents.

Definition 1. Let X, Y, R and S be sets. A *pure open game* $\mathcal{G} = (\Sigma_{\mathcal{G}}, P_{\mathcal{G}}, C_{\mathcal{G}}, E_{\mathcal{G}}) : (X, S) \longrightarrow (Y, R)$ consists of:

- a set $\Sigma_{\mathcal{G}}$, called the set of *strategy profiles* of \mathcal{G} ,
- a function $P_{\mathcal{G}} : \Sigma_{\mathcal{G}} \times X \rightarrow Y$, called the *play function* of \mathcal{G} ,
- a function $C_{\mathcal{G}} : \Sigma_{\mathcal{G}} \times X \times R \rightarrow S$, called the *coutility function* of \mathcal{G} , and
- a function $E_{\mathcal{G}} : X \times (Y \rightarrow R) \rightarrow \mathcal{P}(\Sigma_{\mathcal{G}})$, called the *equilibrium function* of \mathcal{G} . \blacklozenge

As these games are open, they have an interface for interacting with other games. This consists of a set X representing the state/history of the game, a set Y of possible moves, a set R of possible outcomes, and a set S of possible outcomes to feed back to the environment. Open games also have a strategy set $\Sigma_{\mathcal{G}}$ from which we wish to determine the optimal strategy. The play function $P_{\mathcal{G}}$ produces a move based on the state and strategy. The coutility function $C_{\mathcal{G}}$ then determines which outcome is returned to the environment based on the state, strategy and outcome, and the equilibrium function $E_{\mathcal{G}}$ determines which strategies are optimal given the state and utility function. See Example 4 on the next page for an example. The game given there is probabilistic, but as we will see, most of the structure is shared between pure and probabilistic games.

The following fundamental theorem of pure open games allows parallel and sequential composition:

Theorem 2 (Ghani et al. [3]) *The collection of pairs (X, S) of sets X and S , with pure open games $\mathcal{G} : (X, S) \rightarrow (Y, R)$ as morphisms, forms a symmetric monoidal category \mathbf{G}_{Pure} . \blacksquare*

To be precise, in order to satisfy the category axioms on the nose, one needs to quotient by the equivalence relation induced by isomorphism of strategies. We simplify presentation here and in what follows by dealing with representatives directly.

3 Probabilistic Open Games

Our aim is to extend the framework of compositional game theory to also encompass mixed strategies, i.e. games where players' strategies are probability distributions over pure strategies. For a set X , write $\mathcal{D}(X)$ for the set of discrete probability distributions on X , i.e. $\mathcal{D}(X)$ is the collection of functions $\omega : X \rightarrow [0, 1]$ with $\sum_{x \in X} \omega(x) = 1$ whose support $\text{supp}(\omega) = \{x \in X \mid \omega(x) \neq 0\}$ is finite. It is well known that $\mathcal{D} : \text{Set} \rightarrow \text{Set}$ is a monad (see e.g. Jacobs [8] for an overview of probability monads in different categories), and we will make essential use of this structure in the following. The unit of the monad $\eta : X \rightarrow \mathcal{D}X$ maps elements to point distributions, and the multiplication $\mu : \mathcal{D}^2X \rightarrow \mathcal{D}X$ “flattens” a distribution of distributions. Furthermore, \mathcal{D} is a commutative strong monad, meaning that there is a double strength natural transformation $\ell : \mathcal{D}A \times \mathcal{D}B \rightarrow \mathcal{D}(A \times B)$ given by forming the independent joint distribution. Algebras of \mathcal{D} are convex sets, which we think of as sets R equipped with the operation of taking expected values $\mathbb{E} : \mathcal{D}(R) \rightarrow R$. We do not expect all sets involved in a game to support this operation — e.g. the set of moves is typically discrete — but we do expect (and need) the sets of possible outcomes for the games and its environment to do so.

Definition 3. Let X, Y be sets, and R, S be \mathcal{D} -algebras. A *probabilistic open game* $\mathcal{G} = (\Sigma_{\mathcal{G}}, P_{\mathcal{G}}, C_{\mathcal{G}}, E_{\mathcal{G}}) : (X, S) \rightarrow (Y, R)$ consists of:

- a set $\Sigma_{\mathcal{G}}$, called the set of *strategy profiles* of \mathcal{G} ,
- a function $P_{\mathcal{G}} : \Sigma_{\mathcal{G}} \times X \rightarrow Y$, called the *play function* of \mathcal{G} ,
- a function $C_{\mathcal{G}} : \Sigma_{\mathcal{G}} \times X \times R \rightarrow S$, called the *coutility function* of \mathcal{G} , and
- a function $E_{\mathcal{G}} : X \times (Y \rightarrow R) \rightarrow \mathcal{P}(\mathcal{D}(\Sigma_{\mathcal{G}}))$, called the *equilibrium function* of \mathcal{G} . \blacklozenge

In other words, a probabilistic open game consists of the same data as a pure open game, except that the equilibrium function records which *mixed* strategies are “optimal”, instead of just being concerned with pure strategies. Overall, this matches how we usually think of games with mixed strategies: the moves and outcomes of the game stays the same, only the strategies can be probabilistic. The \mathcal{D} -algebra structure of R and S is not needed for this basic definition, but will be used to compose games.

	H	T
H	−1, 1	1, −1
T	1, −1	−1, 1

Figure 1: Utility k of the Matching Pennies game.

Example 4. The Matching Pennies game involves two players trying to win pennies from each other. Each player puts forward one side of a penny, heads or tails. If the faces match then the first player wins the second player’s penny, and if they do not match, the second player instead wins the first player’s penny. This is summarised in Figure 1. We can represent Matching Pennies as a state-free open game

$$\mathcal{M}\mathcal{P} : (\mathbf{1}, \mathbb{R} \times \mathbb{R}) \longrightarrow (\{H, T\} \times \{H, T\}, \mathbb{R} \times \mathbb{R})$$

with utility and coutility taken from $\mathbb{R} \times \mathbb{R}$, and moves $Y \times Y$ where $Y = \{H, T\}$ — each player either plays heads or tails. A pure strategy is simply a move (i.e. the strategy set for the game is $\Sigma_{\mathcal{M}\mathcal{P}} = Y \times Y$), hence both the play and coutility functions $P_{\mathcal{M}\mathcal{P}}$ and $C_{\mathcal{M}\mathcal{P}}$ are particularly simple, given by $P_{\mathcal{M}\mathcal{P}}(c) = c$ and $C_{\mathcal{M}\mathcal{P}}(c, r) = r$ respectively. The equilibrium $E_{\mathcal{M}\mathcal{P}} : (Y \times Y \rightarrow \mathbb{R} \times \mathbb{R}) \rightarrow \mathcal{P}(\mathcal{D}(\Sigma_{\mathcal{M}\mathcal{P}}))$ is defined by $\phi \in E_{\mathcal{M}\mathcal{P}}k$ if and only if

$$\begin{aligned} \phi_1 &\in \arg \max_{\phi'_1 \in \mathcal{D}Y} (\mathbb{E}[\mathcal{D}(\lambda y) \cdot \mathbb{E}[\mathcal{D}(\pi_1 k(y, -))\phi_2]]\phi'_1]) \\ \text{and } \phi_2 &\in \arg \max_{\phi'_2 \in \mathcal{D}Y} (\mathbb{E}[\mathcal{D}(\lambda y') \cdot \mathbb{E}[\mathcal{D}(\pi_2 k(-, y'))\phi_1]]\phi'_2]) \end{aligned}$$

where $\phi_i = \mathcal{D}(\pi_i)\phi$ are the *marginals* of ϕ . We see that both players are trying to maximise their expected payoff, assuming their opponent probabilistically plays according to their fixed strategy. •

4 Probabilistic Open Games Form a Symmetric Monoidal Category

Just like pure open games, probabilistic open games support a wide range of operations: they can be composed in parallel, composed sequentially, conditioned, iterated, and much more. Here we focus on parallel and sequential composition, and prove that these operations make the collection of pairs of sets with probabilistic open games as morphisms a symmetric monoidal category.

4.1 Parallel composition of probabilistic open games

The parallel composition represents two games played simultaneously. Its definition makes crucial use of the fact that the category of \mathcal{D} -algebras has all limits, since it employs products of \mathcal{D} -algebras $R \times R'$ and $S \times S'$.

Definition 5. Let $\mathcal{G} : (X, S) \longrightarrow (Y, R)$ and $\mathcal{G}' : (X', S') \longrightarrow (Y', R')$ be probabilistic open games. We define the *parallel composition* probabilistic open game $\mathcal{G} \otimes \mathcal{G}' : (X \times X', S \times S') \longrightarrow (Y \times Y', R \times R')$ as follows:

- the strategy set is $\Sigma_{\mathcal{G} \otimes \mathcal{G}'} = \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{G}'}$;
- the play function is defined by $P_{\mathcal{G} \otimes \mathcal{G}'}((\sigma, \sigma'), (x, x')) = (P_{\mathcal{G}}(\sigma, x), P_{\mathcal{G}'}(\sigma', x'))$;
- the coutility function is defined by $C_{\mathcal{G} \otimes \mathcal{G}'}((\sigma, \sigma'), (x, x'), (r, r')) = (C_{\mathcal{G}}(\sigma, x, r), C_{\mathcal{G}'}(\sigma', x', r'))$;
- the equilibrium function $E_{\mathcal{G} \otimes \mathcal{G}'} : (X \times X') \times (Y \times Y' \rightarrow R \times R') \rightarrow \mathcal{P}(\mathcal{D}(\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{G}'}))$ is defined by

$$E_{\mathcal{G} \otimes \mathcal{G}'}(x_1, x_2) k = \{ \ell(\phi_1, \phi_2) \mid \phi_1 \in E_{\mathcal{G}} x_1 \mathbb{E}[\mathcal{D}(\pi_1 \circ k) \circ \ell(\eta -, \mathcal{D}(P_{\mathcal{G}'}(-, x_2))\phi_2)] \wedge \\ \phi_2 \in E_{\mathcal{G}'} x_2 \mathbb{E}[\mathcal{D}(\pi_2 \circ k) \circ \ell(\mathcal{D}(P_{\mathcal{G}}(-, x_1))\phi_1), \eta -)] \} \quad \blacklozenge$$

The definition of the strategy set, play function and coutility function coincides with the definition of parallel composition for pure open games, as expected. The equilibrium function of the parallel game is more complicated because of the probabilities involved — note that this makes essential use of the \mathcal{D} -algebra structure on R . Basically, each player is trying to find an equilibrium for the utility function which computes the expected utility for the original utility function $k : Y \times Y' \rightarrow R \times R'$, assuming the other player plays *probabilistically* using their fixed strategy. Note that even though $E_{\mathcal{G} \otimes \mathcal{G}'}(x_1, x_2) k$ is a predicate on $\mathcal{D}(\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{G}'})$, and not on $\mathcal{D}(\Sigma_{\mathcal{G}}) \times \mathcal{D}(\Sigma_{\mathcal{G}'})$, only strategies that arise independently from strategies in $\Sigma_{\mathcal{G}}$ and $\Sigma_{\mathcal{G}'}$ are in the equilibrium. Game-theoretically, this makes sense, as the players are not expected to cooperate, and mathematically, this is crucial for parallel composition to be associative.

Example 6. In Example 4 we showed that the Matching Pennies game can be represented as an open game. We now show that we can build this game as the parallel composition of two identical component “player” games $\mathcal{M}\mathcal{P}_i : (\mathbf{1}, \mathbb{R}) \rightarrow (\{H, T\}, \mathbb{R})$. Strategies are moves $\Sigma_{\mathcal{M}\mathcal{P}_i} = Y = \{H, T\}$ and the play function is given as the identity on strategies. The coutility is given as the second projection returning the utility. Finally the equilibrium function $E_{\mathcal{M}\mathcal{P}_i} : (Y \rightarrow \mathbb{R}) \rightarrow \mathcal{P}(\mathcal{D}(\Sigma_{\mathcal{M}\mathcal{P}_i}))$ is given by

$$\phi \in E_{\mathcal{M}\mathcal{P}_i} k \text{ if } \phi \in \arg \max(\mathbb{E}[\mathcal{D}(k)(-)])$$

i.e. a mixed strategy is optimal if it maximises the expected payoff. The parallel composition of $\mathcal{M}\mathcal{P}_1$ and $\mathcal{M}\mathcal{P}_2$ produces the Matching Pennies game described in Example 4

$$\mathcal{M}\mathcal{P}_1 \otimes \mathcal{M}\mathcal{P}_2 \cong \mathcal{M}\mathcal{P} .$$

The equilibrium function for the composed game states that $\phi \in E_{\mathcal{M}\mathcal{P}_1 \otimes \mathcal{M}\mathcal{P}_2} k$ if

$$\phi_1 \in E_{\mathcal{M}\mathcal{P}_1}(\lambda y . \mathbb{E}[\mathcal{D}(\pi_1 \circ k) \ell(\eta(y), \phi_2)]) \\ \text{and } \phi_2 \in E_{\mathcal{M}\mathcal{P}_2}(\lambda y' . \mathbb{E}[\mathcal{D}(\pi_2 \circ k) \ell(\phi_1, \eta(y'))])$$

where $\phi_i = \mathcal{D}(\pi_i)\phi$ are the marginals of ϕ .

To show that our definition gives the expected results from economic game theory, we now solve this game, i.e. we compute a more concrete description of $E_{\mathcal{M}\mathcal{P}} k$ for the utility function from Figure 1. As Matching Pennies is a symmetric game we focus on the first player’s equilibrium. Expanding the definition of $E_{\mathcal{M}\mathcal{P}_1}$, the condition says

$$\phi_1 \in \arg \max_{\phi'_1 \in \mathcal{D}\Sigma} (\mathbb{E}[\mathcal{D}(\lambda y . \mathbb{E}[\mathcal{D}(\pi_1 \circ k) \ell(\eta(y), \phi_2)])\phi'_1])$$

The vigilant reader might have noticed that the equilibrium condition here is not syntactically the same as the one given in Example 4, but because of the point distributions $\eta(y)$ involved, it is not hard

to see that the expressions are equal. Reducing the terms down and instantiating the utility function from Figure 1, we reach

$$\phi_1 \in \arg \max_{\phi_1 \in D\Sigma} \left(\sum_{r \in R} r \sum_{\{y \in Y \mid \phi_2(y) - \phi_2(\bar{y}) = r\}} \phi_1'(y) \right)$$

As there are only two pure strategies, we can consider both possibilities for ϕ_2 in terms of $\phi_2(H)$ only:

$$\begin{aligned} \phi_2(H) - \phi_2(T) &= \phi_2(H) - (1 - \phi_2(H)) & \phi_2(T) - \phi_2(H) &= (1 - \phi_2(H)) - \phi_2(H) \\ &= 2\phi_2(H) - 1 & &= 1 - 2\phi_2(H) \end{aligned}$$

Rearranging and substituting into the formula, we arrive at the condition

$$\phi_1 \in \arg \max_{\phi'} ((2\phi_2(H) - 1)(2\phi'(H) - 1))$$

and since the game is symmetric we similarly obtain for the second player

$$\phi_2 \in \arg \max_{\phi''} ((2\phi_1(H) - 1)(1 - 2\phi''(H)))$$

leaving three cases to consider:

$$\begin{array}{ll} \text{if } \phi_2(H) = 1/2 \Rightarrow \phi_1(H) \in [0, 1] & \phi_1(H) = 1/2 \Rightarrow \phi_2(H) \in [0, 1] \\ \phi_2(H) < 1/2 \Rightarrow \phi_1(H) = 0 & \phi_1(H) < 1/2 \Rightarrow \phi_2(H) = 1 \\ \phi_2(H) > 1/2 \Rightarrow \phi_1(H) = 1 & \phi_1(H) > 1/2 \Rightarrow \phi_2(H) = 0 \end{array}$$

The only point of stability lies at $\phi_1(H) = \phi_2(H) = 1/2$, since if one player deviates from this strategy the other will return the favour. Hence the only equilibrium is for both players to play both strategies with 50% probability, indeed the standard solution. •

In order to prove associativity of parallel composition, we use a “determinisation” construction that turns probabilistic games into pure games, reminiscent of the abstract categorical formulation of automata determinisation presented e.g. in Silva et al. [12]. This way, we can reuse part of the proof that parallel composition is associative for pure games [3].

Definition 7. Given a probabilistic game $\mathcal{G} : (X, S) \longrightarrow (Y, R)$ with strategy set Σ , we define its *determinisation* pure game $\Delta(\mathcal{G}) : (X, S) \longrightarrow (\mathcal{D}Y, \mathcal{D}R)$ with strategy set $\mathcal{D}\Sigma$ and

- play function $P_{\Delta(\mathcal{G})}(\phi, x) = \mathcal{D}(P_{\mathcal{G}}(-, x))\phi$;
- coutility function $C_{\Delta(\mathcal{G})}(\phi, x, \psi) = \mathbb{E}[\mathcal{D}(C_{\mathcal{G}}(-, x, -))\ell(\phi, \psi)]$; and
- equilibrium function $\phi \in E_{\Delta(\mathcal{G})} x k$ if and only if $\phi \in E_{\mathcal{G}} x (\mathbb{E} \circ k \circ \eta)$. ♦

Using the naturality of η , and that $\mathbb{E} : \mathcal{D}(R) \rightarrow R$ is a \mathcal{D} -algebra, it is easy to see the following way to go between the equilibria of \mathcal{G} and $\Delta(\mathcal{G})$:

Lemma 8 *Let $k : Y \rightarrow R$. Then $\phi \in E_{\Delta(\mathcal{G})} x \mathcal{D}(k)$ if and only if $\phi \in E_{\mathcal{G}} x k$.* ■

In general, it is not the case that the determinisation of a parallel composition is a parallel composition of determinisations — for instance, the type of moves do not even match up, since in general $\mathcal{D}(Y \times Y') \not\cong \mathcal{D}Y \times \mathcal{D}Y'$. To obtain even a lax monoidal map $\Delta(\mathcal{G}) \otimes \Delta(\mathcal{G}') \rightarrow \Delta(\mathcal{G} \otimes \mathcal{G}')$, we need to restrict to utility functions that respect the \mathcal{D} -algebra structure, which for instance Kleisli extensions do. This is formulated in the following lemma.

Lemma 9 Let $\mathcal{G} : (X, S) \longrightarrow (Y, R)$ and $\mathcal{G}' : (X', S') \longrightarrow (Y', R')$ be probabilistic open games. For all $\phi \in \mathcal{D}\Sigma_{\mathcal{G}} \times \mathcal{D}\Sigma_{\mathcal{G}'}$, $x \in X \times X'$, and $k : Y \times Y' \rightarrow \mathcal{D}(R \times R')$, we have

$$\ell(\phi) \in E_{\Delta(\mathcal{G} \otimes \mathcal{G}')} x k^{\#} \text{ iff } \phi \in E_{\Delta(\mathcal{G}) \otimes \Delta(\mathcal{G}')} x ((\mathcal{D}(\pi_1), \mathcal{D}(\pi_2)) \circ k^{\#} \circ \ell)$$

where $k^{\#} = \mu \circ \mathcal{D}(k) : \mathcal{D}(Y \times Y') \rightarrow \mathcal{D}(R \times R')$ is the Kleisli extension of k . ■

We use this lemma to prove the associativity of parallel composition of probabilistic games using the corresponding associativity for pure games.

Theorem 10 Let $\mathcal{G} : (X, S) \longrightarrow (Y, R)$, $\mathcal{G}' : (X', S') \longrightarrow (Y', R')$ and $\mathcal{G}'' : (X'', S'') \longrightarrow (Y'', R'')$ be probabilistic open games. We have $\mathcal{G} \otimes (\mathcal{G}' \otimes \mathcal{G}'') = (\mathcal{G} \otimes \mathcal{G}') \otimes \mathcal{G}''$, up to canonical isomorphisms $A \times (A' \times A'') \cong (A \times A') \times A''$ of the underlying sets involved. ■

4.2 Sequential composition of probabilistic open games

Another fundamental operation to modularly build games is sequential composition. Intuitively, in the sequential composition $\mathcal{G} \circledast \mathcal{H}$ of open games \mathcal{G} and \mathcal{H} , we first play \mathcal{G} , followed by \mathcal{H} . This means the moves of \mathcal{G} are the states of \mathcal{H} , and pure strategies of $\mathcal{G} \circledast \mathcal{H}$ are pairs of pure strategies for \mathcal{G} and \mathcal{H} . A mixed strategy ϕ of the composed game $\mathcal{G} \circledast \mathcal{H}$ is an equilibrium if the marginal distributions are equilibria in \mathcal{G} (relative to the payoff function for \mathcal{G} that we obtain by feeding \mathcal{H} 's coutility back) and \mathcal{H} (relative to the given payoff function of $\mathcal{G} \circledast \mathcal{H}$), respectively. In order to state the latter, we first need to define a ‘‘Kleisli predicate lifting’’ of $E_{\mathcal{H}}(-, k) : Y \rightarrow \mathcal{P}(\mathcal{D}(\Sigma_{\mathcal{H}}))$, since we only get a mixed state in $\mathcal{D}(Y)$ as a result of probabilistically playing the first game using the first mixed strategy.

Definition 11. Let $R : X \rightarrow \mathcal{P}(\mathcal{D}(Y))$. We define $\overline{\mathcal{D}}^{\#}(R) : \mathcal{D}(X) \rightarrow \mathcal{P}(\mathcal{D}(Y))$ by $\overline{\mathcal{D}}^{\#}(R) = \mathcal{P}(\mu_Y) \circ \lambda_{\mathcal{D}(Y)} \circ \mathcal{D}(R)$, where $\lambda : \mathcal{D}\mathcal{P} \rightarrow \mathcal{P}\mathcal{D}$ is the transformation given by

$$\lambda_X(\alpha) = \{\phi \in \mathcal{D}X \mid (\exists \rho \in \mathcal{D}(\subseteq X \times \mathcal{P}X))(\mathcal{D}(\pi_1)\rho = \phi \text{ and } \mathcal{D}(\pi_2)\rho = \alpha)\} . \quad \blacklozenge$$

Concretely, for $\alpha = \sum_i p_i x_i \in \mathcal{D}(X)$, we have

$$\overline{\mathcal{D}}^{\#}(R)(\alpha) = \{\mu(\sum_i \sum_j q_{i,j} \psi_{i,j}) \mid \sum_j q_{i,j} = p_i \text{ and } \psi_{i,j} \in R(x_i)\}$$

where $\sum_i p_i \phi_i$ is the distribution on Y assigning probability $\sum_i p_i \phi_i(y)$ to $y \in Y$. By the abstract definition, we immediately have that $\overline{\mathcal{D}}^{\#}(R \circ f) = \overline{\mathcal{D}}^{\#}(R) \circ \mathcal{D}(f)$ since \mathcal{D} is a functor. We now use this lifting to define the sequential composition of two probabilistic games.

Definition 12. Let $\mathcal{G} : (X, S) \longrightarrow (Y, R)$ and $\mathcal{H} : (Y, R) \longrightarrow (Z, T)$ be probabilistic open games. We define the *sequential composition* probabilistic open game $\mathcal{G} \circledast \mathcal{H} : (X, S) \longrightarrow (Z, T)$ as follows:

- the strategy set is $\Sigma_{\mathcal{G} \circledast \mathcal{H}} = \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}$;
- the play function is defined by $P_{\mathcal{G} \circledast \mathcal{H}}((\sigma_1, \sigma_2), x) = P_{\mathcal{H}}(\sigma_2, P_{\mathcal{G}}(\sigma_1, x))$;
- the coutility function is defined by $C_{\mathcal{G} \circledast \mathcal{H}}((\sigma_1, \sigma_2), x, t) = C_{\mathcal{G}}(\sigma_1, x, C_{\mathcal{H}}(\sigma_2, P_{\mathcal{G}}(\sigma_1, x), t))$;
- the equilibrium function $E_{\mathcal{G} \circledast \mathcal{H}} : X \times (Z \rightarrow T) \rightarrow \mathcal{P}(\mathcal{D}(\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}))$ is defined by

$$E_{\mathcal{G} \circledast \mathcal{H}} x k = \{ \ell(\phi_1, \phi_2) \mid \phi_1 \in E_{\mathcal{G}} x (\lambda y. \mathbb{E}[\mathcal{D}(\lambda \sigma. C_{\mathcal{H}}(\sigma, y, k(P_{\mathcal{H}}(\sigma, y))))\phi_2]) \wedge \phi_2 \in \overline{\mathcal{D}}^{\#}(E_{\mathcal{H}}(-, k))(\mathcal{D}(P_{\mathcal{G}}(-, x))\phi_1) \} \quad \blacklozenge$$




Figure 2: Market Entry game: firms 1 and 2 decide whether to enter (σ_E) or not enter (σ_{NE}) the market.

To see that this definition is meaningful game-theoretically, we model the well-known Market Entry game (Stackelberg [13]) using our framework.

Example 13. The market entry game models two competing firms wishing to enter a new market. If they both enter, the competition between them would be expensive. The situation is depicted in Figure 2. Firm 1 enters first, firm 2 then observes the move made and responds. If one firm enters alone they will reap the rewards, but if both enter they will both suffer. Of course if neither enters then nothing happens. We expect the only subgame perfect equilibrium to be where the first firm enters, and the second firm reverses the first firm’s decision.

We model this as a sequential composition $\mathcal{G}_1 \circledast \mathcal{G}_2$ of two probabilistic open games. The first game $\mathcal{G}_1 : (\mathbf{1}, \mathbb{R} \times \mathbb{R}) \longrightarrow (\{\sigma_E, \sigma_{NE}\}, \mathbb{R} \times \mathbb{R})$ has strategy set $\Sigma_1 = Y = \{\sigma_E, \sigma_{NE}\}$ the set of moves, and the obvious play and cutility functions. Its equilibria are

$$E_1(k : Y \rightarrow \mathbb{R} \times \mathbb{R}) = \arg \max_{\phi \in \mathcal{D}\Sigma_1} \{\mathbb{E}[\mathcal{D}(\pi_1 \circ k)(\phi)]\}$$

The second game $\mathcal{G}_2 : (\{\sigma_E, \sigma_{NE}\}, \mathbb{R} \times \mathbb{R}) \longrightarrow (Y \times Y, \mathbb{R} \times \mathbb{R})$ arises as a “subgame conditioned” game [4, Def. 5] in order to allow the strategies $\Sigma_2 = Y \rightarrow \Sigma_1$ to depend on the move made in \mathcal{G}_1 . The play and cutility functions are given by $P_2(g, x) = (x, g(x))$ and $C_2(g, x, r) = r$. The equilibrium function insists on subgame perfect strategies:

$$\psi \in E_2 y(k : Y \times Y \rightarrow \mathbb{R} \times \mathbb{R}) \quad \text{iff} \quad (\forall y' \in Y) \quad \mathcal{D}(\text{eval}(-, y'))\psi \in \arg \max_{\psi' \in \mathcal{D}\Sigma_1} \{\mathbb{E}[\mathcal{D}(\pi_2 \circ k(y', -))(\psi')]\}$$

where $\text{eval} : (A \rightarrow B) \times A \rightarrow B$ is function evaluation.

The sequential composition $\mathcal{G}_1 \circledast \mathcal{G}_2 : (\mathbf{1}, \mathbb{R} \times \mathbb{R}) \longrightarrow (Y \times Y, \mathbb{R} \times \mathbb{R})$ has as strategies pairs of strategies from each round $\Sigma_{\mathcal{G}_1 \circledast \mathcal{G}_2} = \Sigma_1 \times \Sigma_2$. For mixed strategies $\phi \in \mathcal{D}\Sigma_1$ and $\psi \in \mathcal{D}\Sigma_2$, we have $\ell(\phi, \psi) \in E_{\mathcal{G}_1 \circledast \mathcal{G}_2}(k : Y \times Y \rightarrow \mathbb{R} \times \mathbb{R})$ if and only if

$$\begin{aligned} \phi &\in E_1(\lambda y. \mathbb{E}[\mathcal{D}(\lambda f. k(y, f(y)))\psi]) \quad \text{and} \\ \psi &\in \overline{\mathcal{D}}^\#(E_2(-, k))\phi = \mathcal{D}(E_2(\sigma_E, k)) = \mathcal{D}(E_2(\sigma_{NE}, k)) \end{aligned}$$

where the second condition has been simplified since $E_2(y, k)$ is independent of y . For the utility function k from Figure 2, we further see that in fact $E_2(\sigma_E, k) = E_2(\sigma_{NE}, k) = \{1 \cdot \text{swap}\}$ where $\text{swap} : Y \rightarrow Y$ is the function which swaps σ_E and σ_{NE} . Hence for $\ell(\phi, \psi) \in E_{\mathcal{G}_1 \circledast \mathcal{G}_2} k$ we must have $\psi = 1 \cdot \text{swap}$ which in turns forces $\phi = 1 \cdot \sigma_E$ — the expected (non-mixed) subgame perfect equilibria. Reflecting on the argument, we see that our reasoning is an instance of backward induction (see e.g. Leyton-Brown and Shoham [10, §4.4]). It is interesting, but currently not clear to us, what the $\overline{\mathcal{D}}^\#(-)$ construction does in general when the second game has not been conditioned to respond to the moves of the first game. •

It is important to note that the distributive law $\lambda : \mathcal{D}\mathcal{P} \rightarrow \mathcal{P}\mathcal{D}$ used in Definition 11 is *not* a distributive law between monads, because no such law exists (Zwart and Marsden [14]). In particular, λ does not preserve the monad structure of \mathcal{D} , for instance $\lambda_X \circ \eta_{\mathcal{D}X} \neq \eta_{\mathcal{D}X}^{\mathcal{P}}$. It is however a distributive law between functors (even of a functor over the monad \mathcal{P} , and also over \mathcal{P}^{op} , although we do not make use of this fact), which will be important for us for proving associativity of sequential composition.

Fact 14 *The transformation $\lambda : \mathcal{D}\mathcal{P} \rightarrow \mathcal{P}\mathcal{D}$ is a distributive law between functors, i.e. it is natural.* ■

For a proof see Kupke, Kurz and Venema [9]. Using the naturality of λ , we can show that if $R : X \rightarrow \mathcal{P}\mathcal{D}Y$ and $f : \mathcal{D}Y \rightarrow \mathcal{D}Y'$, then $\mathcal{P}(f) \circ \overline{\mathcal{D}}^{\#}(R) = \overline{\mathcal{D}}^{\#}(\mathcal{P}(f) \circ R)$. This, with f being a marginal $\mathcal{D}(\pi)$, is one of the key steps to prove associativity of composition.

Theorem 15 *Let $\mathcal{G} : (X, S) \rightarrow (X', S')$, $\mathcal{G}' : (X', S') \rightarrow (X'', S'')$ and $\mathcal{G}'' : (X'', S'') \rightarrow (Y, R)$ be probabilistic open games. We have $\mathcal{G} \circ (\mathcal{G}' \circ \mathcal{G}'') = (\mathcal{G} \circ \mathcal{G}') \circ \mathcal{G}''$, up to the canonical isomorphism $\Sigma_{\mathcal{G}} \times (\Sigma_{\mathcal{G}'} \times \Sigma_{\mathcal{G}''}) \cong (\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{G}'}) \times \Sigma_{\mathcal{G}''}$ of strategy sets.* ■

4.3 A symmetric monoidal category

We have now assembled most of pieces needed to show that probabilistic open games are the morphisms of a monoidal category: missing are unit and identity games.

For each set X and \mathcal{D} -algebra S , we define a probabilistic open game $\mathcal{I}\mathcal{D}_{(X,S)} : (X, S) \rightarrow (X, S)$ with strategy set $\Sigma_{\mathcal{I}\mathcal{D}_{(X,S)}} = \mathbf{1}$, play function $P_{\mathcal{I}\mathcal{D}_{(X,S)}}(\sigma, x) = x$, coutility function $C_{\mathcal{I}\mathcal{D}_{(X,S)}}(\sigma, x, s) = s$, and equilibrium function $E_{\mathcal{I}\mathcal{D}_{(X,S)}} xk = \mathbf{1}$, i.e. every (trivial) strategy is an equilibrium.

Lemma 16 *There is a category \mathbf{G}_{Prob} , where objects are pairs (X, S) of a set X and a \mathcal{D} -algebra S , and morphisms are probabilistic open games. Composition is given by sequential composition $\mathcal{G} \circ \mathcal{H} = \mathcal{H} \circ \mathcal{G}$, and the identity on (X, S) is $\mathcal{I}\mathcal{D}_{(X,S)}$.* ■

Similarly, we define a trivial game $\mathcal{I} : (\mathbf{1}, \mathbf{1}) \rightarrow (\mathbf{1}, \mathbf{1})$ with strategy set $\Sigma_{\mathcal{I}} = \mathbf{1}$, the only possible play and coutility functions, and equilibrium function $E_{\mathcal{I}} xk = \mathbf{1}$, i.e. every strategy is again an equilibrium.

Lemma 17 *The game \mathcal{I} is the unit for parallel composition. Furthermore, the operation which maps (X, S) and (X', S') to $(X \times X', S \times S')$, and games \mathcal{G} and \mathcal{G}' to $\mathcal{G} \otimes \mathcal{G}'$, defines a bifunctor $\otimes : \mathbf{G}_{\text{Prob}} \times \mathbf{G}_{\text{Prob}} \rightarrow \mathbf{G}_{\text{Prob}}$.* ■

Observing that \mathbf{G}_{Prob} also has a symmetry (inherited from $\text{Set} \times (\mathcal{D}\text{-Alg})^{\text{op}}$), we have now proved the following:

Theorem 18 *The collection of pairs (X, S) of a set X and a \mathcal{D} -algebra S , with probabilistic open games $\mathcal{G} : (X, S) \rightarrow (Y, R)$ as morphisms, forms a symmetric monoidal category \mathbf{G}_{Prob} .* ■

5 Relating pure and probabilistic games

We now construct a category where probabilistic open games are the objects, by defining a notion of morphism between games. In light of Theorem 18, these morphisms are 2-cells in a monoidal double category of games (cf. Hedges [7]). The construction works similarly for pure games. We then use the resulting categorical structure to relate pure and probabilistic games in the form of an adjunction between the categories.

As noticed by Ghani et al. [3], the definition of pure open games can be given more compactly by employing the language of lenses [2]. A lens $(v, u) : (X, S) \rightarrow (Y, R)$ between pairs of sets (X, S)

and (Y, R) is given by a two functions $v : X \rightarrow Y$ (“view”) and $u : X \times R \rightarrow S$ (“update”). Hence the play and coutility functions of a game $\mathcal{G} : (X, S) \rightarrow (Y, R)$ can equivalently be described as a family of lenses $(P_{\mathcal{G}}(\sigma, -), C_{\mathcal{G}}(\sigma, -, -)) : (X, S) \rightarrow (Y, R)$ indexed by strategies $\sigma \in \Sigma_{\mathcal{G}}$. Further, the data involved in the equilibrium function can be described by a “global element” lens $(\mathbf{1}, \mathbf{1}) \rightarrow (X, S)$ and a “global co-element” lens $(Y, R) \rightarrow (\mathbf{1}, \mathbf{1})$. As a result, most reasoning about open games can be done diagrammatically using that lenses also compose: given $(v, u) : (X, S) \rightarrow (Y', R')$ and $(v', u') : (Y', R') \rightarrow (Y, R)$, we can construct a lens $(X, S) \rightarrow (Y, R)$ by $(v' \circ v : X \rightarrow Y, (x, y) \mapsto u(x, u'(v(x), y))) : X \times Y \rightarrow R$.

There is an identity-on-objects functor $\iota(-, -) : \text{Set} \times \text{Set}^{\text{op}} \rightarrow \text{Lens}$ that maps a pair of functions $(f : X \rightarrow Y, g : R \rightarrow S)$ to a lens $\iota(f, g) : (X, S) \rightarrow (Y, R)$ with f as first component and $g \circ \pi_2 : X \times R \rightarrow S$ as second component.

Definition 19. Let $\mathcal{G} : (X, S) \rightarrow (Y, R)$ and $\mathcal{G}' : (X', S') \rightarrow (Y', R')$ be pure (probabilistic) open games. A *morphism of pure (probabilistic) games* $\mathcal{G} \rightarrow \mathcal{G}'$ consists of functions

$$(f_P : X \rightarrow X', f_C : S' \rightarrow S) \quad (g_P : Y \rightarrow Y', g_C : R' \rightarrow R)$$

and $h : \Sigma_{\mathcal{G}} \rightarrow \Sigma_{\mathcal{G}'}$, such that the following diagram of lenses commutes for each $\sigma \in \Sigma_{\mathcal{G}}$

$$\begin{array}{ccc} (X, S) & \xrightarrow{\iota(f_P, f_C)} & (X', S') \\ (P_{\mathcal{G}}(\sigma), C_{\mathcal{G}}(\sigma)) \downarrow & & \downarrow (P_{\mathcal{G}'}(h(\sigma)), C_{\mathcal{G}'}(h(\sigma))) \\ (Y, R) & \xrightarrow{\iota(g_P, g_C)} & (Y', R') \end{array}$$

and, for every $x \in X$ and $k : Y' \rightarrow R'$, we have that $\sigma \in E_{\mathcal{G}} x (g_C \circ k \circ g_P)$ implies

- $h(\sigma) \in E_{\mathcal{G}'} (f_P(x)) k$ for pure games,
- $\mathcal{D}(h)(\sigma) \in E_{\mathcal{G}'} (f_P(x)) k$ for probabilistic games.

We write $\text{Game}_{\text{Prob}}$ and $\text{Game}_{\text{Pure}}$ for the categories of probabilistic and pure open games, respectively, where the morphisms are defined as above. \blacklozenge

This is a generalisation of the definition of morphism of state-free games used in our paper on iterated open games [4], but different from the notion of morphism employed by Hedges [7], which fails to make the determinisation operation Δ from Definition 7 a functor. As there are currently a number of viable notions of morphisms of games (even of lenses), we consider this empirical evidence important for what an appropriate notion of morphism for games ought to be. For the rest of this section, let $\text{Game}'_{\text{Pure}}$ be the category $\text{Game}_{\text{Pure}}$, except that utility and coutility sets are additionally endowed with \mathcal{D} -algebra structure.

Proposition 20 A variant of determinisation Δ' mapping a probabilistic game $\mathcal{G} : (X, S) \rightarrow (Y, R)$ to a pure game $\Delta'(\mathcal{G}) : (\mathcal{D}X, \mathcal{D}S) \rightarrow (\mathcal{D}Y, \mathcal{D}R)$ (using the double strength ℓ , and $\overline{\mathcal{D}}^{\#}(-)$), still with strategy set $\Sigma_{\Delta'(\mathcal{G})} = \mathcal{D}(\Sigma_{\mathcal{G}})$, extends to a functor $\Delta' : \text{Game}_{\text{Prob}} \rightarrow \text{Game}'_{\text{Pure}}$. \blacksquare

Determinisation Δ itself is a functor if restricted to games whose coutility preserves the \mathcal{D} -algebra structure in a certain sense. One might hope that one of these functors might have a left or a right adjoint, but this is too much to ask, since it would imply in turn that \mathcal{D} has both a left and a right adjoint. However, we show that the canonical way to embed a pure game as a probabilistic game has a right adjoint.

Theorem 21 Let $\Theta : \text{Game}'_{\text{Pure}} \rightarrow \text{Game}_{\text{Prob}}$ be the functor that acts as the identity on the strategy set and the lens structure, with $E_{\Theta(\mathcal{G})} xk = \{\eta(\sigma) \mid \sigma \in E_{\mathcal{G}} xk\}$. Then

$$\begin{array}{ccc} & \xleftarrow{\Theta} & \\ \text{Game}_{\text{Prob}} & \perp & \text{Game}'_{\text{Pure}} \\ & \xrightarrow{\Psi} & \end{array}$$

where $\Psi : \text{Game}_{\text{Prob}} \rightarrow \text{Game}'_{\text{Pure}}$ similarly acts as the identity on the strategy set and the lens structure, with $E_{\Psi(\mathcal{H})} xk = \{\sigma \mid \eta(\sigma) \in E_{\mathcal{H}} xk\}$. ■

6 Conclusions and Future Work

We have presented a framework for compositional game theory which encompasses also mixed strategies, and shown that it is closed under parallel and sequential composition, and shown that it can adequately model common games such as Matching Pennies (where mixed strategies are crucial) and the Market Entry Game. We also defined a notion of morphism between games, and showed that it gives rise to a category of games that we that can be useful for reasoning, e.g. by employing adjunctions between pure and probabilistic games.

Several challenges remain. While we have accurately captured mixed strategy Nash equilibria — a fundamental solution concept in game theory — it remains to be seen if this framework can exploit the non-independent distributions that arise naturally in it to capture also *correlated equilibria* or perhaps even *evolutionary stable strategies*. Finally, we remark that most of our proofs do not use any particular properties of the commutative monad \mathcal{D} . We think this can be used to uniformly model other “effectful” game-theoretic phenomena such as e.g. quitting games using the exceptions monad.

Acknowledgements We thank the participants at the Third Symposium on Compositional Structures (SYCO3) in Oxford, March 2019 for discussions and comments.

References

- [1] Martín Escardó & Paulo Oliva (2010): *Sequential games and optimal strategies*. *Proceedings of the Royal Society of London A* 467, pp. 1519–1545, doi:10.1098/rspa.2010.0471.
- [2] J. Nathan Foster, Michael B. Greenwald, Jonathan T. Moore, Benjamin C. Pierce & Alan Schmitt (2007): *Combinators for Bidirectional Tree Transformations: A Linguistic Approach to the View-update Problem*. *ACM Trans. Program. Lang. Syst.* 29(3), doi:10.1145/1232420.1232424.
- [3] Neil Ghani, Jules Hedges, Viktor Winschel & Philipp Zahn (2018): *Compositional Game Theory*. In: *Logic in Computer Science*, LICS '18, pp. 472–481, doi:10.1145/3209108.3209165.
- [4] Neil Ghani, Clemens Kupke, Alasdair Lambert & Fredrik Nordvall Forsberg (2018): *A compositional treatment of iterated open games*. *Theor. Comput. Sci.* 741, pp. 48–57, doi:10.1016/j.tcs.2018.05.026.
- [5] Michele Giry (1982): *A categorical approach to probability theory*. In: *Categorical aspects of topology and analysis*, Springer, pp. 68–85, doi:10.1007/BFb0092872.
- [6] Jules Hedges (2016): *Towards compositional game theory*. Ph.D. thesis, Queen Mary University London.
- [7] Jules Hedges (2018): *Morphisms of Open Games*. *Electr. Notes Theor. Comput. Sci.* 341, pp. 151–177, doi:10.1016/j.entcs.2018.11.008.
- [8] Bart Jacobs (2018): *From probability monads to commutative effectuses*. *J. Log. Algebr. Meth. Program.* 94, pp. 200–237, doi:10.1016/j.jlamp.2016.11.006.

- [9] Clemens Kupke, Alexander Kurz & Yde Venema (2012): *Completeness for the coalgebraic cover modality*. *Logical Methods in Computer Science* 8(3), doi:10.2168/LMCS-8(3:2)2012.
- [10] Kevin Leyton-Brown & Yoav Shoham (2008): *Essentials of Game Theory: A Concise Multidisciplinary Introduction*. Morgan and Claypool, doi:10.2200/S00108ED1V01Y200802AIM003.
- [11] John Nash (1951): *Non-Cooperative Games*. *Ann. Math.* 54(2), pp. 286–295, doi:10.2307/1969529.
- [12] Alexandra Silva, Filippo Bonchi, Marcello Bonsangue & Jan Rutten (2013): *Generalizing determinization from automata to coalgebras*. *Logical Methods in Computer Science* 9(1), doi:10.2168/LMCS-9(1:9)2013.
- [13] Heinrich von Stackelberg (2011): *Market structure and equilibrium*. Springer, doi:10.1007/978-3-642-12586-7. Translation from the German language edition “Marktform und Gleichgewicht” (1934).
- [14] Maaïke Zwart & Dan Marsden (2019): *Don’t Try This at Home: No-Go Theorems for Distributive Laws*. In: *Logic in Computer Science*, LICS ’19, pp. 1–13, doi:10.1109/LICS.2019.8785707.