# Bayesian open games 

Joe Bolt, Jules Hedges ${ }^{1,2,3}$, and Philipp Zahn ${ }^{2,3}$<br>${ }^{1}$ Department of Computer and Information Sciences, University of Strathclyde, 26 Richmond Street, Glasgow, G11XH, U.K.<br>${ }^{2} 20$ squares, https://20squares.xyz<br>${ }^{3}$ CyberCat Institute, https://cybercat.institute

This paper generalises the treatment of compositional game theory as introduced by Ghani et al. in 2018, where games are modelled as morphisms of a symmetric monoidal category. From an economic modelling perspective, the notion of a game in the work by Ghani et al. is not expressive enough for many applications. This includes stochastic environments, stochastic choices by players, as well as incomplete information regarding the game being played. The current paper addresses these three issues all at once.

## 1 Introduction

In [13] the first compositional treatment of economic game theory was introduced. Following the literature on categorical open systems [7], open games are modelled as morphisms of a symmetric monoidal category.

A distinctive and non-obvious feature of this approach is that the Nash equilibrium condition [30], one of the central concepts in classical game theory to analyse rational behaviour of agents (cf. [12, Chapter 1.2] and [33, Chapter 2]), is itself compositional.

While an important first step, the treatment in [13] has two severe limitations:

1. Games are deterministic and as a consequence, there are no chance elements in the games and players have to choose deterministically.
2. Players have complete information about all relevant data of the game such as payoffs, number of players etc.

Many interesting strategic situations feature chance elements. Poker is one example - already discussed in the ground-breaking work of von Neumann and Morgenstern [31]. In an economic context, the environment also often is non-deterministic. Two competing companies face uncertain demand, exchange rates, lawsuits etc.

More subtle but also important is that players may need and may want to randomise their actions. There are well known situations like Matching Pennies (see, for instance, [12, p. 16]) where playing deterministically means being 'beaten' all the time. Conceptually, from a game theory perspective, this means that there are games where equilibria do not exist when players are limited to deterministic strategies (known as 'pure strategies') whereas they do exist when players can choose stochastically (known as 'mixed strategies').

Lastly, it is a crude approximation to assume that players have complete information. Examples abound. A used-car dealer knows how good the car is that he is trying to sell to you. You may not know. Banks sitting on toxic assets know how little value they actually have. The government trying to buy these assets in order to save the financial system from collapse may not know. An agent bidding in an auction may not know how many other bidders he competes with. In most situations incomplete information is the norm and not the exception.

[^0]The above limitations restrict the applicability of compositional game theory to economic phenomena. And it restricts its usefulness for economists. After all, classical game theory already deals with these complications.

In this paper we provide a generalisation of open games which solves the three problems above in one go. We adapt the core definition of compositional game theory such that the environment can be stochastic and players can also choose in a non-deterministic fashion. Doing so, we also introduce a way to deal with incomplete information. Essentially we are lifting the same 'trick', which has been introduced in classical game theory to deal with games of incomplete information by John Harsanyi [17-19], to compositional game theory.

Harsanyi argued that instead of dealing with games of incomplete information directly, which poses formidable conceptual problems, we can transform such games into games of imperfect information by introducing the notion of (game) types. Players have access to probability distributions characterising these games as well as partial access to this information. For instance, in an auction a player may know how much he values the good to be auctioned. However, he may not know how other agents value the good. Assuming that players update information according to Bayes' rule and adapting the equilibrium notion of Nash, to what is called Bayesian Nash equilibrium, game theorists can work with such interactions no differently to how they deal with chance elements in Poker. Thus by transforming the problem, Harsanyi essentially opened the path to using tools that were more or less already introduced by von Neumann and Morgenstern [31].

We are applying the same strategy. By introducing stochastic environments and adapting the equilibrium notion from Nash to Bayesian Nash we show that our compositional framework captures exactly Bayesian Games and thus allows to deal with stochastic environments as well as with situations of incomplete information. This distinguishes our work also from [15] which addresses the issue of deterministic players in isolation.

### 1.1 Technical introduction

Contrary at least to our own initial beliefs, addressing these issues requires some significant adaptions of open games as defined in [13].

The recent understanding of open games has been based on lenses, which consist of a pair of functions $X \rightarrow Y$ and $X \times R \rightarrow S$ packaged into a single morphism $(X, S) \rightarrow(Y, R)$ of a category. Here, the function $X \rightarrow Y$ is the play function, which plays out a given strategy by taking an initial state to a final state of the open game. The function $X \times R \rightarrow S$, known as the coplay function or coutility function, is more subtle: It 'backpropagates' payoffs into the past, given an initial state. This operation is 'counterfactual', and the composition of lenses (which is not trivial to define, nor is obvious to see is associative) intertwines ordinary forward and counterfactual (or 'teleological') information flow.

An open game can then be viewed as a family of lenses indexed by a set of strategy profiles, together with another component describing which strategy profiles are Nash equilibria in a given context. A context for an open game consists of an initial state $(X)$ and a function from final states to payoffs $(Y \rightarrow R)$. Contexts turn out also to be intimately connected to lenses, and indeed this was the initial hint that viewing open games in terms of lenses is a deep idea rather than a coincidence.

To someone trained in thinking about processes with side effects, it is entirely natural to begin by inserting a (finite support) probability monad $D$, and take the components of the lenses to be Kleisli morphisms $X \rightarrow D(Y)$ and $X \times R \rightarrow D(S)$, or equivalently to use lenses over the category of sets and (finite support) stochastic functions. This allows the strategies of an open game to describe probabilistic behaviours, which are known as behavioural strategies in game theory. Unfortunately this doesn't work: In order to prove that lenses form a category (i.e. are associative and unital) it is necessary that the forwards maps $X \rightarrow Y$ are homomorphisms of copying, and in the category of stochastic processes this characterises those processes that are actually deterministic.

Fortunately this problem has already been solved in the theory of lenses, although the solution is far from obvious. We use the existential lenses or coend lenses as developed by Riley [36]. This means we replace the pair or functions $X \rightarrow D(Y)$ and $X \times R \rightarrow D(S)$ with three things: A choice of set $A$, a function $X \rightarrow D(A \times Y)$ and a function $A \times R \rightarrow D(S)$. Moreover a certain equivalence relation needs to be imposed, and this is precisely given by the following coend [27] (one of the
universal constructions of category theory):

$$
\int^{A}(X \rightarrow D(A \times Y)) \times(A \times R \rightarrow D(S))
$$

The proofs in this paper make heavy use of a diagrammatic language for existential lenses developed in [36].

The second question is what should be considered a context of a Bayesian open game, i.e. a replacement for the pair $X \times(Y \rightarrow R)$. There is an existing characterisation of these contexts in terms of deterministic lenses, namely as a 'state' lens $(1,1) \rightarrow(X, S)$ and a 'costate' lens $(Y, R) \rightarrow(1,1)$. However this turns out to be a red herring: in Section 3.4 we show that generalising from this causes the category of open games to fail to be monoidal in an unexpected way.

It turns out that the appropriate notion of context consists of three things: a set $\Theta$ of unobservable states, a joint distribution on $\Theta \times X$ (i.e. an element of $D(\Theta \times X)$ ) and a function $\Theta \times Y \rightarrow D(R)$. Again we need to impose a certain equivalence relation, which again turns out to be precisely a coend. Remarkably this is equivalent to a state in the category of double lenses, i.e. lenses over the category of lenses. This brings an unexpected theoretical unity to Bayesian open games, and means that the graphical language of [36] can be used throughout.

## 2 Concrete open games

We begin with a self-contained introduction to deterministic open games. In essence, we will introduce the necessary machinery so that we can represent simple classical games with diagrams as depicted below.


This diagram displays an interaction between two agents, A and B. Player A moves first; player B observes the choice by A and then moves afterwards. Both moves are consumed by an environment $c_{k}$ which provides the payoff for both players.

To get to a full understanding of this diagram (and deterministic open games), in this section we introduce several building blocks. Roughly, they can be classified in two kinds.

First, we need a way to architecture the information flow. As we will see in Sections 2.1 to 2.3 , lenses play a crucial role by providing us with a categorical structure on which to build open games.

Second, we need to flesh out the internals of the boxes in the diagram. Specifically, how does strategic reasoning actually take place? Central here is the notion of an agent who makes observations and chooses a course of action. As we will see in Section 2.5 and Section 2.6, the key insight is to model an agent as choosing against a context which comprises how the environment reacts to an agent's choices. The context is also the glue that keeps the outside information flow and the internal reasoning together.

Once we have introduced all the relevant parts, we will come back to the example above.
Note: The exposition in this section slightly deviates from previous work. We believe this eases the way for the generalisations to come in Section 3 and thereafter.

### 2.1 Lenses

The history of mathematical lenses is complicated, involving many independent discoveries and fresh starts across numerous areas of mathematics and computer science [2, 5, 9, 26, 32, 34]. An in-depth description of this history can be found at [25]. We use lenses to describe the flow of
information through a game. A lens for a given game describes which players have access to what information when making a strategic decision, and also how information about players' strategic decisions is ultimately fed into the outcome function for the game. For example, it may specify an order of play, or whether two players are playing in parallel, or even whether some players are privy to certain information in the environment that other players are not.

In general, lenses can be thought of as processes that perform some computation and then propagate some resulting feedback from the environment backwards through a system of which they are a part. In particular, this means that lenses have both covariant and contravariant components. The covariant component carries out the initial computation and the contravariant component propagates the resulting feedback back through the system. Crucially, lenses are also compositional in the sense that they admit both sequential and parallel composition and, consequently, form a symmetric monoidal category.

The lenses used in this paper are direct descendants of the lenses of database theory. Given some data $x$ of type $X$ we may want to view some part of it $y$ of type $Y$. This is encapsulated by a view function $v: X \rightarrow Y$. From this 'close-up' view of the database we may want to edit the database by updating $y$. Given an update of the view $y$ we then need to know how this update propagates to an update of the original data $x$. That is, given initial data $x$ and an updated view $y^{\prime}: Y$, we should specify some updated $x^{\prime}: X$ given by some update function $u: X \times Y \rightarrow X$. The pair $(v, u)$ is a lens with type $X \rightarrow Y$. The connection to our previous abstract definition of lenses is as follows:

- The covariant computation associated with the lens is the view function $v: X \rightarrow Y$,
- the resulting feedback from the environment is the update made to the subdatabase returned by the view function, and
- this feedback is propagated back to the whole database via the update function $u: X \times Y \rightarrow$ $X$.

Abstracting away from databases, there is no reason to demand that the feedback generated by the environment will have the same type as the output of the lens computation. Similarly, we may be interested in cases where the update function is not-so-literally an 'update' function, but merely a function that propagates some kind of feedback back through the system. As such, the lenses we will be using will have types of the form $(X, S) \rightarrow(Y, R)$ where the covariant component of the lens is of type $X \rightarrow Y$ and the contravariant component is of type $X \times R \rightarrow S$.

In game theory, we can regard players as 'lenses that care about the feedback they receive from the environment'. In a game with sequential play, players make some play (computation), receive some utility (feedback) from the outcome function, and then pass some feedback to earlier players in the game (their outcome function given the moves that the later players chose). Moreover, given that lenses admit parallel composition as well as sequential composition, we obtain a nuanced notion of information flow in a game.

In the next subsections we describe a symmetric monoidal category of concrete lenses. 'Concrete' here refers to the fact that the view and update functions are functions in Set. We then come to the core of this section, the definition of a concrete open game.

### 2.2 The category of concrete lenses

Definition 2.2.1 (Concrete lens). Let $X, S, Y$ and $R$ be sets. A concrete lens $l:(X, S) \rightarrow(Y, R)$ is a pair of functions $\left(l_{v}: X \rightarrow Y, l_{u}: X \times R \rightarrow S\right)$.

As a trivial first example, there is an obvious mapping that takes a morphism of $\boldsymbol{\operatorname { S e t }} \times \boldsymbol{S e t}^{\mathrm{op}}$ and returns a concrete lens.

Example 2.2.2. Let $f: X \rightarrow Y$ and $g: R \rightarrow S$. Define a concrete lens $\langle f, g\rangle:(X, S) \rightarrow(Y, R)$ by

$$
\begin{gathered}
\langle f, g\rangle_{v}=f \\
\langle f, g\rangle_{u}(x, r)=g(r)
\end{gathered}
$$

Definition 2.2.3 (Sequential composition of concrete lenses). Let $l:(X, S) \rightarrow(Y, R)$ and $t$ : $(Y, R) \rightarrow(Z, Q)$ be concrete lenses. The sequential composite $t \circ l:(X, S) \rightarrow(Z, Q)$ is given by $\left((t \circ l)_{v}: X \rightarrow Z,(t \circ l)_{u}: X \times Q \rightarrow S\right)$ where

$$
(t \circ l)_{v}=t_{v} \circ l_{v}
$$

and $(t \circ l)_{u}$ is given by

$$
X \times Q \xrightarrow{\Delta_{X} \times \mathrm{id}_{X}} X \times X \times Q \xrightarrow{\mathrm{id}_{X} \times l_{v} \times \mathrm{id}_{Q}} X \times Y \times Q \xrightarrow{\mathrm{id}_{X} \times t_{u}} X \times R \xrightarrow{l_{u}} S .
$$

As a string diagram $(t \circ l)_{u}$ is given by


Lemma 2.2.4 (Sequential composition of concrete lenses is associative). Suppose we have concrete lenses

$$
(X, S) \xrightarrow{l}(Y, R) \xrightarrow{m}(Z, Q) \xrightarrow{n}(W, T)
$$

Then $n \circ(m \circ l)=(n \circ m) \circ l$.
Theorem 2.2.5 (Concrete lenses form a category). There is a category CL with pairs of sets as objects and concrete lenses as morphisms.

### 2.3 The monoidal structure of concrete lenses

Definition 2.3.1 (Tensor composition of concrete lenses). Let $l_{1}:\left(X_{1}, S_{1}\right) \rightarrow\left(Y_{1}, R_{1}\right)$ and $l_{2}$ : $\left(X_{2}, S_{2}\right) \rightarrow\left(Y_{2}, R_{2}\right)$ be concrete lenses. The tensor composition $l_{1} \otimes l_{2}:\left(X_{1} \times X_{2}, S_{1} \times S_{2}\right) \rightarrow$ $\left(Y_{1} \times Y_{2}, R_{1} \times R_{2}\right)$ is given by $\left(\left(l_{1} \otimes l_{2}\right)_{v},\left(l_{1} \otimes l_{2}\right)_{u}\right)$ where

$$
\left(l_{1} \otimes l_{2}\right)_{v}=l_{1_{v}} \times l_{2_{v}}
$$

and $\left(l_{1} \otimes l_{2}\right)_{u}$ is given by

$$
X_{1} \times X_{2} \times R_{1} \times R_{2} \xrightarrow{\cong} X_{1} \times R_{1} \times X_{2} \times R_{2} \xrightarrow{l_{1_{u}} \times l_{2_{u}}} S_{1} \times S_{2}
$$

In a diagram, $\left(l_{1} \otimes l_{2}\right)_{u}$ is


Lemma 2.3.2. $\otimes$ is a functor.
Theorem 2.3.3. There is a symmetric monoidal category $\mathbf{C L}$ where the objects are pairs of sets and the morphisms are concrete lenses. Sequential composition and the monoidal tensor are as in the above definitions. The monoidal unit is $I=(\{*\},\{*\})$.

The following observations about states and effects in CL will be useful in the remainder of this section.

Lemma 2.3.4. $\mathbf{C L}(I,(X, S)) \cong X$.
Proof. This is easily seen, as a state $l \in \mathbf{C L}(I,(X, S))$ is given by a pair

$$
(s:\{\star\} \rightarrow X, e:\{\star\} \times S \rightarrow\{\star\}) .
$$

Lemma 2.3.5. $\mathbf{C L}((Y, R), I) \cong(Y \rightarrow R)$
Proof. An effect $l \in \mathbf{C L}((Y, R), I)$ is given by a pair

$$
(v: Y \rightarrow\{\star\}, u: Y \times\{\star\} \rightarrow R) .
$$

### 2.4 Concrete open games

Now we have the necessary prerequisites in place to introduce the notion of a concrete open game. A concrete open game consists of a set of strategy profiles; a family of concrete lenses indexed by the set of strategy profiles; and a best-response function.

Definition 2.4.1 (Concrete open game). Let $X, S, Y$, and $R$ be sets. A concrete open game $\mathcal{G}:(X, S) \rightarrow(Y, R)$ is given by

1. A set of strategy profiles $\Sigma$;
2. A play function $\mathrm{P}: \Sigma \rightarrow \mathbf{C L}((X, S),(Y, R))$; and
3. A best-response function $\mathrm{B}: X \times(Y \rightarrow R) \rightarrow \operatorname{Rel}(\Sigma)$.

Here $\operatorname{Rel}(\Sigma)=\mathcal{P}(\Sigma \times \Sigma)$ is the set of binary relations on $\Sigma$. In this paper we will usually specify a relation in terms of its forward images $\Sigma \rightarrow \mathcal{P}(\Sigma)$.

The type $X$ is the type of observations made by the game; the type $Y$ is the type of actions that can be chosen; the type $R$ is the type of outcomes; and the type $S$ is the type of co-outcomes. Of the four types associated with a concrete open game, the type $S$ is the most mysterious. Succinctly, its purpose is to relay information about outcomes to games acting earlier. In a sequential composite $\mathcal{H} \circ \mathcal{G}$ of open games (we will define sequential composition of concrete open games shortly), the co-outcome type of $\mathcal{H}$ is also the outcome type of $\mathcal{G}$. We think of $\mathcal{H}$ as receiving some outcome which is then acted upon by the contravariant component of a concrete lens given by $\mathcal{H}$ 's play function before being passed back to $\mathcal{G}$ as $\mathcal{G}$ 's outcome.

The best-response function of an open game is an abstraction from the utility functions of classical game theory. Recall that a Nash equilibrium for a normal-form game is a strategy profile in which no player has incentive to unilaterally deviate. We can instead think of a relation on the set of strategy profiles for a normal-form game where strategy profiles $\sigma$ and $\tau$ are related if $\tau$ is the result of players unilaterally deviating from $\sigma$ to their most profitable unilateral deviation. Nash equilibria are then the fixed points of this relation. For convenience, in the definition of a concrete open game we work directly with a best-response relation rather than preference relations.

The play function takes a strategy as argument and returns a concrete lens that describes an open play of the game $\mathcal{G}$ ('open' here means 'lacking a particular observation and outcome function' and is explained in the next paragraph). To justify this interpretation, recall that a concrete lens $l:(X, S) \rightarrow(Y, R)$ consists of $v: X \rightarrow Y$ and $u: X \times R \rightarrow S$. The view function $v$ describes how a game decides on an action given an observation (similar to how strategies for sequential games work). The update function $u$ describes precisely how games relay information about outcomes to other games acting earlier.

As the name suggests, concrete open games are open to their environment. The appropriate notion of a context for a concrete open game is given in the following definition. A concrete open game together with a context can be thought of as a full description of a game.

Definition 2.4.2. Let $\mathcal{G}:(X, S) \rightarrow(Y, R)$. A history for $\mathcal{G}$ is an element $x$ of $X$, an outcome function for $\mathcal{G}$ is a function $k: Y \rightarrow R$, and a context for $\mathcal{G}$ is a pair $(x, k): X \times(Y \rightarrow R)$.

We are now in a position to justify the type of the best-response function. The best response functions takes a context as argument, and a context is precisely the information required for resolving the 'openness' of a concrete open game. Given a context, the best response function then returns the set of best deviations from a strategy profile $\sigma$.

We represent a concrete open game $\mathcal{G}:(X, S) \rightarrow(Y, R)$ using the diagram


This diagrammatic notation emphasises the point that information flows both covariantly through $\mathcal{G}$ from observations to actions, and contravariantly through $\mathcal{G}$ from outcomes to cooutcomes. These diagrams constitute a bona fide diagrammatic calculus for the category of concrete open games defined in the remainder, as detailed in [21].

Notation 2.4.3. String diagrams in the category of open games will always be drawn with arrowheads on wires, whilst string diagrams in the ambient category will always be drawn without arrowheads.

Atomic concrete open games are an important class of concrete open games, and are the basic components out of which more complex games are constructed. Whilst concrete open games can, in general, represent aggregates of agents responding to each other (in a way that will be made precise in 2.7 and 2.8), atomic concrete open games describe games in which there is no strategic interaction. Examples are simple computations in which no decisions are made whatsoever, and single agents that are sensitive only to a given context.

Definition 2.4.4 (Atomic concrete open game). A concrete open game $a:(X, S) \rightarrow(Y, R)$ is atomic if

1. $\Sigma_{a} \subseteq \mathbf{C L}((X, S),(Y, R))$;
2. For all $l \in \Sigma_{a}, \mathrm{P}_{\mathcal{G}}(l)=l$; and
3. For all contexts $c: X \times(Y \rightarrow R), \mathrm{B}_{a}(c): \Sigma_{a} \rightarrow \mathcal{P}\left(\Sigma_{a}\right)$ is constant.

We sometimes refer to an atomic concrete open game simply as an atom.
Note that an atom $a:(X, S) \rightarrow(Y, R)$ is fully determined by a subset $\Sigma_{a} \subseteq \mathbf{C L}((X, S),(Y, R))$ and a selection function ${ }^{1} \varepsilon: X \times(Y \rightarrow R) \rightarrow \mathcal{P}\left(\Sigma_{a}\right)$.

Given $f: X \rightarrow Y$ and $g: R \rightarrow S$, as in CL, the pair $(f, g) \in \boldsymbol{\operatorname { S e t }} \times \boldsymbol{S e t}^{\mathrm{op}}$ can be represented as a concrete open game. We refer to such games as computations, as no strategic choice is being made.

Example 2.4.5 (Computation). Let $f: X \rightarrow Y$ and $g: R \rightarrow S$. The atom $\langle f, g\rangle:(X, S) \rightarrow(Y, R)$ is given by

1. $\Sigma=\{\langle f, g\rangle\}$; and
2. For all $c: X \times(Y \rightarrow R), \varepsilon(c)=\{\langle f, g\rangle\}$.

Similar to CL, the following computations will turn out to be the underlying structural maps for the symmetric monoidal category of concrete open games.

[^1]Definition 2.4.6 (Structural computations). Define identity, associator, swaps, and left/right unitor computations to be the atomic concrete open games given by

$$
\begin{aligned}
\operatorname{id}_{(X, S)} & =\left\langle\operatorname{id}_{X}, \mathrm{id}_{S}\right\rangle \\
\alpha_{X \otimes(Y \otimes Z), A \otimes(B \otimes C)} & =\left\langle\alpha_{X, Y, Z}, \alpha_{A, B, C}^{-1}\right\rangle \\
s_{(X, A),(Y, B)} & =\left\langle s_{(X, Y)}, s_{(A, B)}^{-1}\right\rangle \\
\rho_{(X, Y)} & =\left\langle\rho_{X}, \rho_{Y}^{-1}\right\rangle \\
\lambda_{(X, Y)} & =\left\langle\lambda_{X}, \lambda_{Y}^{-1}\right\rangle
\end{aligned}
$$

where the Set functions on the right-hand side of the equalities are the obvious Set isomorphisms.
Counit games are an interesting class of atoms that reverse the direction of information flow in a concrete open game.

Definition 2.4.7 (Counit). Let $f: X \rightarrow S$. Define an atomic concrete open game $c_{f}:(X, S) \rightarrow$ $(\{\star\},\{\star\})$ by

1. $\Sigma_{c_{f}}=\{\langle!, f\rangle\}$; and
2. For all $c: X \times(\{\star\} \rightarrow\{\star\}), \varepsilon(c)=\Sigma_{c_{f}}$.

We are being slightly relaxed with notation here as the update function for $c_{f}$ has type $X \times$ $\{\star\} \rightarrow S$ while $f$ has type $X \rightarrow S$. We represent $c_{f}$ as follows.


### 2.5 Agents

So far we have only seen open games for which the set of strategies is a singleton, describing games with no strategic decisions. Our first examples of a concrete open game with non-trivial strategy set are agents. These can be used to represent the utility maximising agents of traditional game theory or, more generally, to represent players trying to influence the outcome of a game.

Definition 2.5.1 (Agent). An agent $\mathcal{A}:(X,\{\star\}) \rightarrow(Y, R)$ is an atom whose set of strategies is $\Sigma=\mathbf{C L}((X,\{\star\}),(Y, R))$.

Recall that a concrete lens $l: \mathbf{C L}((X,\{\star\}),(Y, R))$ is a pair $(v: X \rightarrow Y, u: X \times R \rightarrow\{\star\})$ and, hence, is uniquely determined by a function of type $X \rightarrow Y$. Consequently, a strategy for an agent specifies how an agent map chooses an action of type $Y$ given an observation of type $X$. Given a context $c: X \times(Y \rightarrow R), \mathrm{B}_{\mathcal{A}}(c)$ picks out the set of strategies $\mathcal{A}$ considers acceptable in the context $c$. Agents are represented diagrammatically by


We can specialise the definition above to model the utility-maximising agents of traditional game theory.

Example 2.5.2 (Utility maximising agent). The utility maximising agent $\mathcal{A}:(X,\{\star\}) \rightarrow(Y, \mathbb{R})$ is given by

$$
\varepsilon(x, k)=\{\sigma: X \rightarrow Y \mid \sigma(x) \in \arg \max (k)\} .
$$

There are other decision criteria one could use. For instance, MinMax and regret minimisation would be candidates. We could also consider models from behavioural game theory such as prospect theory. The only (very weak) requirement is that the decision criterion can be described by a selection function [23, 24]. In this paper we focus on utility-maximising agents as a simplification and in order to match traditional game theory.

### 2.6 Best-response with concrete lenses

Recall from 2.3.4 and 2.3.5 that $\mathbf{C L}(I,(X, S)) \cong X$ and that $\mathbf{C L}(I,(Y, R)) \cong Y \rightarrow R$. Using these facts we can rephrase the type of best response for a concrete open game $\mathcal{G}:(X, S) \rightarrow(Y, R)$ as

$$
\mathrm{B}_{\mathcal{G}}: \mathbf{C L}(I,(X, S)) \times \mathbf{C L}((Y, R), I) \rightarrow \operatorname{Rel}\left(\Sigma_{\mathcal{G}}\right)
$$

This formulation allows for a concise and natural definition of sequential composition for concrete open games where it would otherwise seem $a d$ hoc. To make matters clear, we write $x$ when talking about elements of $X$ and $x^{\star}$ when talking about concrete lenses with type $\mathbf{C L}(I,(X, S))$. Similarly, we write $k: Y \rightarrow R$ when talking about functions in Set and we write $k_{\star}$ when talking about effects in $\mathbf{C L}((Y, R), I)$.

### 2.7 Sequential composition of concrete open games

In this section we specify how to define the sequential composite $\mathcal{H} \circ \mathcal{G}:(X, S) \rightarrow(Z, Q)$ of two concrete open games $\mathcal{G}:(X, S) \rightarrow(Y, R)$ and $\mathcal{H}:(Y, R) \rightarrow(Z, Q)$.

We imagine that this composition really is sequential in a straightforward way. $\mathcal{G}$ is 'played out' according to some strategy $\sigma \in \Sigma_{\mathcal{G}}$ and then $\mathcal{H}$ is 'played out' according to some $\tau \in \Sigma_{\mathcal{H}}$. A choice of $(\sigma, \tau) \in \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}$ therefore determines an open play of $\mathcal{G}$ and $\mathcal{H}$ played in sequence, and so we take $\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}$ to be the set of strategy profiles of $\mathcal{H} \circ \mathcal{G}$.

The play function of the sequential composite is defined straightforwardly using the sequential composition of concrete lenses defined in 2.2.3.

Defining best response for a sequential composite is a bit more delicate and, for explanatory purposes, we make use of the informal notion of a local context for a subgame. Given a context $c=(x: X, k: Z \rightarrow Q)$ and a strategy $(\sigma, \tau)$ for $\mathcal{H} \circ \mathcal{G}$, the best-response relation of $\mathcal{H} \circ \mathcal{G}$ is specified by calling the best-response function of $\mathcal{G}$ with a modified context corresponding to how $c$ 'appears' to $\mathcal{G}$ when $\mathcal{H}$ plays according to $\tau$ and, similarly, calling the best-response function of $\mathcal{H}$ with a modified context corresponding to how $c$ 'appears' to $\mathcal{H}$ when $\mathcal{G}$ plays according to $\sigma$. In practice we define these 'local contexts' in the obvious way that type checks, but this is because the work has already been done in carefully choosing the correct definitions.

Definition 2.7.1 (Sequential composition for concrete open games). Let $\mathcal{G}=\left(\Sigma_{\mathcal{G}}, \mathrm{P}_{\mathcal{G}}, \mathrm{B}_{\mathcal{G}}\right)$ : $(X, S) \rightarrow(Y, R)$ and $\mathcal{H}=\left(\Sigma_{\mathcal{H}}, \mathrm{P}_{\mathcal{H}}, \mathrm{B}_{\mathcal{H}}\right):(Y, R) \rightarrow(Z, Q)$ be concrete open games. Define

1. $\Sigma_{\mathcal{H} \circ \mathcal{G}}=\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}$,
2. $\mathrm{P}_{\mathcal{H} \circ \mathcal{G}}(\sigma, \tau)=\mathrm{P}_{\mathcal{H}}(\tau) \circ \mathrm{P}_{\mathcal{G}}(\sigma)$ (where $\circ$ composition is in $\mathbf{C L}$ ), and
3. $\mathrm{B}_{\mathcal{H} \circ \mathcal{G}}\left(x^{\star}, k_{\star}\right)(\sigma, \tau)=\mathrm{B}_{\mathcal{G}}\left(x^{\star}, k_{\star} \circ \mathrm{P}_{\mathcal{H}}(\tau)\right) \sigma \times \mathrm{B}_{\mathcal{H}}\left(\mathrm{P}_{\mathcal{G}}(\sigma) \circ x^{\star}, k_{\star}\right)(\tau)$.

We represent $\mathcal{H} \circ \mathcal{G}$ with the diagram


### 2.8 Tensor composition for concrete open games

The tensor composition of open games represents simultaneous play. Given concrete open games $\mathcal{G}:\left(X_{1}, S_{1}\right) \rightarrow\left(Y_{1}, R_{1}\right)$ and $\mathcal{H}:\left(X_{2}, S_{2}\right) \rightarrow\left(Y_{2}, R_{2}\right)$, the strategy set for $\mathcal{G} \otimes \mathcal{H}$ is $\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}} ;$ we make use of the tensor composition in $\mathbf{C L}$ in defining the play function; and the best-response function is given by modifying the context $c$ to give local contexts for $\mathcal{G}$ and $\mathcal{H}$.

Definition 2.8.1 (Local contexts for tensor composition). Define the left local tensor context operator

$$
\mathcal{L}:\left(X^{\prime} \times\left(X^{\prime} \rightarrow Y^{\prime}\right) \times\left(Y \times Y^{\prime} \rightarrow R \times R^{\prime}\right)\right) \rightarrow(Y \rightarrow R)
$$

by

$$
\mathcal{L}\left(x^{\prime}, p^{\prime}, k\right)(y)=\pi_{1} \circ k\left(y, p^{\prime}\left(x^{\prime}\right)\right) .
$$

As a diagram, $\mathcal{L}\left(x^{\prime}, p^{\prime}, l\right)$ is the function


Similarly, define the right local tensor context operator

$$
\mathcal{R}:\left(X \times(X \rightarrow Y) \times\left(Y \times Y^{\prime} \rightarrow R \times R^{\prime}\right)\right) \rightarrow\left(Y^{\prime} \rightarrow R^{\prime}\right)
$$

by

$$
\mathcal{R}(x, p, k)\left(y^{\prime}\right)=\pi_{2} \circ k(p(x), y)
$$

As a diagram,


Suppose we have concrete open games $\mathcal{G}:(X, S) \rightarrow(Y, R)$ and $\mathcal{H}:\left(X^{\prime}, S^{\prime}\right) \rightarrow\left(Y^{\prime}, R^{\prime}\right)$ and we wish to combine them to create some game $\mathcal{G} \otimes \mathcal{H}:\left(X \times X^{\prime}, S \times S^{\prime}\right) \rightarrow\left(Y \times Y^{\prime}, R \times R^{\prime}\right)$. Consider the left context operator $\mathcal{L}$ acting on some triple ( $x^{\prime}, p^{\prime}, k$ ). If $k$ is an outcome function for the game $\mathcal{G} \otimes \mathcal{H}$ and $\mathcal{H}$ observes $x^{\prime}$ and plays according to the function $p^{\prime}$, then $\mathcal{L}\left(x^{\prime}, p^{\prime}, k\right)$ is the 'apparent' outcome function for $\mathcal{G}$. Similarly, $\mathcal{R}(x, p, k)$ is the 'apparent' outcome function for $\mathcal{H}$ when $\mathcal{G}$ observes $x$ and plays according to $p$. With this in mind, we define tensor composition for concrete open games as follows.

Definition 2.8.2 (Tensor composition of concrete open games). Let $\mathcal{G}:(X, S) \rightarrow(Y, R)$ and $\mathcal{H}:\left(X^{\prime}, S^{\prime}\right) \rightarrow\left(Y^{\prime}, R^{\prime}\right)$ be concrete open games. Define

$$
\mathcal{G} \otimes \mathcal{H}:\left(X \times X^{\prime}, S \times S^{\prime}\right) \rightarrow\left(Y \times Y^{\prime}, R \times R^{\prime}\right)
$$

by

1. $\Sigma_{\mathcal{G} \otimes \mathcal{H}}=\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}$,
2. $\mathrm{P}_{\mathcal{G} \otimes \mathcal{H}}(\sigma, \tau)=\mathrm{P}_{\mathcal{G}}(\sigma) \otimes \mathrm{P}_{\mathcal{H}}(\tau)$ (in $\left.\mathbf{C L}\right)$,
3. $\mathrm{B}_{\mathcal{G} \otimes \mathcal{H}}:\left(\left(X \times X^{\prime}\right) \times\left(Y \times Y^{\prime} \rightarrow R \times R^{\prime}\right)\right) \rightarrow \operatorname{Rel}\left(\Sigma_{\mathcal{G} \otimes \mathcal{H}}\right)$ is given by

$$
\begin{aligned}
\mathrm{B}_{\mathcal{G} \otimes \mathcal{H}}\left(\left(x, x^{\prime}\right)^{\star}, k_{\star}\right)(\sigma, \tau)= & \mathrm{B}_{\mathcal{G}}\left(x^{\star}, \mathcal{L}\left(x^{\prime},\left(\mathrm{P}_{\mathcal{H}}(\tau)\right)_{v}, k\right)_{\star}\right)(\sigma) \\
& \times \mathrm{B}_{\mathcal{H}}\left(x^{\prime \star}, \mathcal{R}\left(x,\left(\mathrm{P}_{\mathcal{G}}(\sigma)\right)_{v}, k\right)_{\star}\right)(\tau)
\end{aligned}
$$

$\mathcal{G} \otimes \mathcal{H}$ is represented by the diagram


### 2.9 Equivalence of open games

One subtlety remains before we can define the category of concrete open games. We aim to define a category with pairs of sets as objects and morphisms given by concrete open games. If carried out naïvely, this runs into the problem that strategy sets which should be identical are merely isomorphic. For instance, the strategy set of $\mathcal{K} \circ(\mathcal{H} \circ \mathcal{G})$ is $\left(\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}\right) \times \Sigma_{\mathcal{K}}$ whilst the strategy set of $(\mathcal{K} \circ \mathcal{H}) \circ \mathcal{G}$ is $\Sigma_{\mathcal{G}} \times\left(\Sigma_{H} \times \Sigma_{\mathcal{K}}\right)$. In order for concrete open games to form a category, we must first take an appropriate quotient.

There are several different reasonable choices of quotient. Since this is an orthogonal consideration to this paper's topic, we choose the most straightforward, which is to identify open games that have a compatible isomorphism between their sets of strategies. Other choices that can be made are bisimulations [3] and surjections [8]. Alternatively, instead of taking a quotient, open games can be considered as the 1-cells of a bicategory [22].

Definition 2.9.1. Let $\mathcal{G}, \mathcal{H}:(X, S) \rightarrow(Y, R)$ be concrete open games. An isomorphism $\alpha: \mathcal{G} \rightarrow \mathcal{H}$ is given by a bijection $\alpha: \Sigma_{\mathcal{G}} \rightarrow \Sigma_{\mathcal{H}}$ such that

1. $\mathrm{P}_{\mathcal{G}}(\sigma)=\mathrm{P}_{\mathcal{H}}(\alpha(\sigma))$ for all $\sigma \in \Sigma_{\mathcal{G}}$
2. For all $\sigma, \sigma^{\prime} \in \Sigma_{\mathcal{G}}$ and $c \in X \times(Y \rightarrow R),\left(\sigma, \sigma^{\prime}\right) \in \mathrm{B}_{\mathcal{G}}(c)$ iff $\left(\alpha(\sigma), \alpha\left(\sigma^{\prime}\right)\right) \in \mathrm{B}_{\mathcal{H}}(c)$

Definition 2.9.2. Let $\mathcal{G}, \mathcal{H}:(X, S) \rightarrow(Y, R)$ be concrete open games. $\mathcal{G}$ and $\mathcal{H}$ are equivalent, written $\mathcal{G} \sim \mathcal{H}$, if there exists an isomorphism $\alpha: \mathcal{G} \rightarrow \mathcal{H}$. We write [G] for the equivalence class of $\mathcal{G}$ under this relation. We also say that the isomorphism $\alpha$ witnesses the equivalence between $\mathcal{G}$ and $\mathcal{H}$ and write $\mathcal{G} \stackrel{\alpha}{\sim} \mathcal{H}$.

The following results demonstrate that sequential and tensor composition of concrete open respects equivalence of concrete open games.

Lemma 2.9.3. Let $\mathcal{G}, \mathcal{G}^{\prime}:(X, S) \rightarrow(Y, R)$ and $\mathcal{H}, \mathcal{H}^{\prime}:(Y, R) \rightarrow(Z, Q)$ be concrete open games. If $\mathcal{G} \sim \mathcal{G}^{\prime}$ and $\mathcal{H} \sim \mathcal{H}^{\prime}$, then $\mathcal{H} \circ \mathcal{G} \sim \mathcal{H}^{\prime} \circ \mathcal{G}^{\prime}$.

Proof. Suppose $\mathcal{G} \stackrel{\alpha}{\sim} \mathcal{H}$ and $\mathcal{G}^{\prime} \stackrel{\beta}{\sim} \mathcal{H}^{\prime}$. Then $\alpha \times \beta: \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{G}} \rightarrow \Sigma_{\mathcal{G}^{\prime}} \times \Sigma_{\mathcal{H}^{\prime}}$ given by

$$
(\alpha \times \beta)(\sigma, \tau)=(\alpha(\sigma), \beta(\tau))
$$

is such that $\mathcal{H} \circ \mathcal{G} \stackrel{\alpha \times \beta}{\sim} \mathcal{H}^{\prime} \circ \mathcal{G}^{\prime}$.
Lemma 2.9.4. Let $\mathcal{G}, \mathcal{H}:(X, S) \rightarrow(Y, R)$ and $\mathcal{G}^{\prime}, \mathcal{H}^{\prime}:\left(X^{\prime}, S^{\prime}\right) \rightarrow\left(Y^{\prime}, R^{\prime}\right)$ be concrete open games. If $\mathcal{G} \sim \mathcal{H}$ and $\mathcal{G}^{\prime} \sim \mathcal{H}^{\prime}$, then $\mathcal{G} \otimes \mathcal{G}^{\prime} \sim \mathcal{H} \otimes \mathcal{H}^{\prime}$.

Proof. If $\mathcal{G} \stackrel{\alpha}{\sim} \mathcal{H}$ and $\mathcal{G}^{\prime} \stackrel{\beta}{\sim} \mathcal{H}^{\prime}$, then $\mathcal{G} \otimes \mathcal{G}^{\prime} \stackrel{\alpha \times \beta}{\sim} \mathcal{H} \otimes \mathcal{H}^{\prime}$ as in the previous lemma.

### 2.10 The category of concrete open games

We are now finally in a position to show that concrete open games form a symmetric monoidal category.

Notation 2.10.1. In string diagrams we refer to a play function applied to a strategy simply by the strategy. For example, $\sigma$ may refer to $\mathrm{P}_{\mathcal{G}}(\sigma)$. In practice this does not lead to ambiguity because proofs and definitions proceed by assigning fixed strategies to particular open games. This notational convention allows for less cluttered string diagrams.

Lemma 2.10.2. Sequential composition of concrete open games is associative up to equivalence.
The identity morphism $(X, S) \rightarrow(X, S)$ is given by the computation $\left\langle\mathrm{id}_{X}, \mathrm{id}_{S}\right\rangle$.
Lemma 2.10.3. Let $\mathcal{G}:(X, S) \rightarrow(Y, R)$. Then $[\mathcal{G}]=\left[\mathcal{G} \circ\left\langle\operatorname{id}_{X}, \operatorname{id}_{S}\right\rangle\right]=\left[\left\langle\operatorname{id}_{Y}, \operatorname{id}_{R}\right\rangle \circ \mathcal{G}\right]$.
Corollary 2.10.4. There is a category ConGame with pairs of sets as object and equivalence classes of concrete open games as morphisms.

We now move on to proving that ConGame is symmetric monoidal.
Lemma 2.10.5. $\otimes$ : ConGame $\times$ ConGame $\rightarrow$ ConGame is a functor .
Lemma 2.10.6. The associator in ConGame is natural.
We include the proof of this lemma specifically because it will be important later.
Proof. Let $\mathcal{G}_{i}:\left(X_{i}, S_{i}\right) \rightarrow\left(Y_{i}, R_{i}\right)$ be open games where $i \in\{1,2,3\}$. We need to show that $\alpha \circ\left(\mathcal{G}_{1} \otimes\left(\mathcal{G}_{2} \otimes \mathcal{G}_{3}\right)\right) \sim\left(\left(\mathcal{G}_{1} \otimes \mathcal{G}_{2}\right) \otimes \mathcal{G}_{3}\right) \circ \alpha$. Define $\beta:\left(\Sigma_{\mathcal{G}_{1}} \times\left(\Sigma_{\mathcal{G}_{2}} \times \Sigma_{\mathcal{G}_{3}}\right)\right) \times\{\star\} \rightarrow\{\star\} \times\left(\left(\Sigma_{\mathcal{G}_{1}} \times\right.\right.$ $\left.\left.\Sigma_{\mathcal{G}_{2}}\right) \times \Sigma_{\mathcal{G}_{3}}\right)$ by $\beta((\sigma,(\tau, \mu)), \star)=(\star,((\sigma, \tau), \mu))$.

As CL is symmetric monoidal, we have that

$$
\alpha \circ \mathrm{P}_{\mathcal{G}_{1} \otimes\left(\mathcal{G}_{2} \otimes \mathcal{G}_{3}\right)}(\sigma,(\tau, \mu))=\mathrm{P}_{\left(\mathcal{G}_{1} \otimes \mathcal{G}_{2}\right) \otimes \mathcal{G}_{2}}((\sigma, \tau), \mu) \circ \alpha .
$$

Let $x_{i} \in X_{i}$ and $k:\left(Y_{1} \times Y_{2} \times Y_{3}\right) \rightarrow\left(R_{1} \times R_{2} \times R_{3}\right)$. We will show that the local contexts for $\mathcal{G}_{1}, \mathcal{G}_{2}$, and $\mathcal{G}_{3}$ are the same in both $\mathcal{G}_{1} \otimes\left(\mathcal{G}_{2} \otimes \mathcal{G}_{2}\right)$ and $\left(\mathcal{G}_{1} \otimes \mathcal{G}_{2}\right) \otimes \mathcal{G}_{3}$. First we consider $\mathcal{G}_{1}$. Let $k^{\prime}:=\mathcal{L}\left(x_{3}, \mathrm{P}_{\mathcal{G}_{3}} \mu, k\right)$. Then,

$$
k_{\mathcal{G}_{1}}:=\mathcal{L}\left(\left(x_{2}, x_{3}\right), \mathrm{P}_{\mathcal{G}_{2} \otimes \mathcal{G}_{3}}(\tau, \mu), k\right)
$$



Similar arguments hold for $\mathcal{G}_{2}$ and $\mathcal{G}_{3}$, showing that

$$
k_{\mathcal{G}_{2}}:=\mathcal{R}\left(x_{1}, \mathrm{P}_{\mathcal{G}_{1}}(\sigma), \mathcal{L}\left(x_{3}, \mathrm{P}_{\mathcal{G}_{3}}(\mu), k\right)\right)=\mathcal{L}\left(x_{3}, \mathrm{P}_{\mathcal{G}_{3}}(\mu), \mathcal{R}\left(x_{1}, \mathrm{P}_{\mathcal{G}_{1}}(\sigma), k\right)\right)
$$

and

$$
k_{\mathcal{G}_{3}}:=\mathcal{R}\left(\left(x_{1}, x_{2}\right), \mathrm{P}_{\mathcal{G}_{1} \otimes \mathcal{G}_{2}}(\sigma, \tau), k\right)=\mathcal{R}\left(x_{2}, \mathrm{P}_{\mathcal{G}_{2}}(\tau), \mathcal{R}\left(x_{1}, \mathrm{P}_{\mathcal{G}_{1}}(\sigma), k\right)\right)
$$

Then

$$
\begin{aligned}
&\left((\star,(\sigma,(\tau, \mu))),\left(\star,\left(\sigma^{\prime},\left(\tau^{\prime}, \mu^{\prime}\right)\right)\right)\right) \\
& \Longleftrightarrow \in \mathrm{B}_{\alpha \circ}\left(\mathcal{G}_{1} \otimes\left(\mathcal{G}_{2} \otimes \mathcal{G}_{3}\right)\right) \\
&\left(\left(x_{1},\left(x_{2}, x_{3}\right)\right), k\right) \\
& \Longleftrightarrow\left(\sigma, \sigma^{\prime}\right) \in \mathrm{B}_{\mathcal{G}_{1}}\left(x_{1}, k_{\mathcal{G}_{1}}\right) \text { and }\left(\tau, \tau^{\prime}\right) \in \mathrm{B}_{\mathcal{G}_{2}}\left(x_{2}, k_{\mathcal{G}_{2}}\right) \text { and }\left(\mu, \mu^{\prime}\right) \in \mathrm{B}_{\mathcal{G}_{3}}\left(x_{3}, k_{\mathcal{G}_{3}}\right) \\
&\left.\left((((\sigma, \tau), \mu), \star),\left(\left(\left(\sigma^{\prime}, \tau^{\prime}\right), \mu^{\prime}\right), \star\right)\right)\right) \in \mathrm{B}_{\left(\left(\mathcal{G}_{1} \otimes \mathcal{G}_{2}\right) \otimes \mathcal{G}_{3}\right) \circ \alpha}\left(\left(\left(x_{1}, x_{2}\right), x_{3}\right), k\right)
\end{aligned}
$$

The above lemma relies on the fact that the monoidal tensor in Set is cartesian. In particular we needed that bipartite states $s: I \rightarrow S_{1} \otimes S_{2}$ in Set (i.e. elements of $S_{1} \times S_{2}$ ) correspond to pairs of states $\left(s_{1}: I \rightarrow S_{1}, s_{2}: I \rightarrow S_{2}\right)$. In an arbitrary monoidal category, it need not be the case that for all states $s: I \rightarrow S_{1} \otimes S_{2}$ there exist states $s_{1}: I \rightarrow S_{1}$ and $s_{2}: I \rightarrow S_{2}$ such that


This poses a significant barrier to generalising concrete open games to monoidal categories where the monoidal tensor is not cartesian, and Section 3 addresses this problem.

Lemma 2.10.7. The structural computations $\lambda, \rho$, and $s$ are natural in ConGame.
Theorem 2.10.8. ConGame is symmetric monoidal.

### 2.11 Encoding functions as games

Recall that, given functions $f: X \rightarrow Y$ and $g: R \rightarrow S$, there is a computation of concrete open games $\langle f, g\rangle:(X, S) \rightarrow(Y, R)$. In fact, this operation is functorial.

Lemma 2.11.1 ([20]). Define $F:$ Set $\times$ Set $^{\text {op }} \rightarrow$ ConGame by $F(X, S)=(X, S)$ and $F(f:$ $X \rightarrow Y, g: R \rightarrow S)=\langle f, g\rangle$. Then $F$ is a faithful monoidal functor.

We also incorporate computations directly into the diagrammatic calculus for concrete open games, representing the computation $\langle f, g\rangle:(X, S) \rightarrow(Y, R)$ by


Two particularly useful examples of this notation are the covariant and contravariant copying computations $\left\langle\Delta_{X}, \mathrm{id}_{1}\right\rangle:(X, 1) \rightarrow(X \times X, 1)$ and $\left\langle\mathrm{id}_{1}, \Delta_{R}\right\rangle:(1, R \times R) \rightarrow(1, R)$ which are represented by
 and

respectively.

### 2.12 Game theory with concrete open games

In this section we give some examples of games modelled using concrete open games. We will be light on details, aiming to simply demonstrate some of the expressive power of concrete open games. We direct the reader to [20] for more details.

### 2.12.1 Bimatrix games

Bimatrix games are simply two-player normal-form games, the most well-known example of which is likely the prisoner's dilemma. We assume the set of actions available to each player is finite for simplicity.

Definition 2.12.2. A bimatrix game consists of

1. Finite set of actions $A$ and $B$; and
2. An outcome function $k: A \times B \rightarrow \mathbb{R}^{2}$.

A bimatrix game $\mathcal{G}=(A, B, k)$ is represented by the concrete open game

where $\mathcal{A}$ and $\mathcal{B}$ are utility maximising agents and $c_{k}$ is the counit game associated with $k$. In diagrammatic form, the structure of the game is made clear. Players $\mathcal{A}$ and $\mathcal{B}$ make independent choices from $A$ and $B$ respectively which are then used to generate two real numbers as outcomes. Bimatrix games may not have a Nash equilibrium in pure strategies, but in cases that do have Nash equilibria, they appear as fixed points of the best-response function B : $A \times B \rightarrow \mathcal{P}(A \times B)$ of the above concrete open game, i.e. strategy profiles $(a, b)$ satisfying $(a, b) \in \mathrm{B}(a, b)$.

### 2.12.3 Two-player sequential game

A two-player sequential game is defined by the same data as a bimatrix game (sets $A$ and $B$ and a function $k: A \times B \rightarrow \mathbb{R}^{2}$ ), but we allow the second player to observe the first player's move before making a choice, so strategies for the second player are functions $A \rightarrow B$. This is represented by the concrete open game

where $\mathcal{A}$ and $\mathcal{B}$ are utility maximising agents and $c_{k}$ is the counit game associated with $k$.
Crucially, the fixed points of the concrete open game are not subgame perfect Nash equilibria, but rather plain old Nash equilibria. It is also possible to define a concrete open game that captures subgame perfect equilibria, but this requires an additional operator defined in [14].

### 2.12.4 Normal-form games

Let $\Gamma=\left(N,\left(S_{i}\right)_{i=1}^{N},\left(u_{i}\right)_{i=1}^{N}\right)$ be a normal-form game for $N$ players. $S_{i}$ denotes the set of strategies available to player $i$. The function $u_{i}$ maps a strategy tuple for all players, $\prod_{i=1}^{N} S_{i}$, to player i's payoff, in $\mathbb{R}$. Define $k: \prod_{i=1}^{N} S_{i} \rightarrow \mathbb{R}^{N}$ by $s=\left(s_{1}, \cdots, s_{N}\right) \mapsto\left(u_{1}(s), \cdots, u_{N}(s)\right)$. We can model this normal-form game using the concrete open game

$$
c_{k} \circ\left(\bigotimes_{i=1}^{N} \mathcal{A}_{i}\right)
$$

where $\mathcal{A}_{i}: I \rightarrow\left(S_{i}, \mathbb{R}\right)$ is the utility maximising agent. The fixed points of this game's best-response relation are then the pure-strategy Nash equilibria of the normal-form game.

## 3 General open games

The notion of open game we introduced in the section before can emulate some standard games such as the prisoner's dilemma. On the other hand, classical game theory has a much wider reach. It can model situations with which a concrete open game cannot deal. This involves stochastic environments, probabilistic choices by players, and incomplete information.

In this section, we will significantly generalise the notion of an open game, to make room for these three extensions (and beyond). The first order of business, to make progress in this direction, is to generalise concrete lenses.

### 3.1 Generalising concrete lenses

In the proof of Lemma 2.2.4 we made use of the fact that every Set function is a comonoid homomorphism for the copy/delete comonoid. Recall that a morphism is a comonoid homomorphism if it can be 'moved through' the comonoid structure.


If Set is replaced with some arbitrary symmetric monoidal category $\mathcal{C}$ and the copy/delete comonoid is replaced with some arbitrary comonoid in $\mathcal{C}$, sequential composition of lenses, as defined in Definition 2.2.3, may not be associative. This presents a substantive problem - there exist categories relevant to game theory in which sequential composition of concrete lenses is not associative. Of particular interest is the Kleisli category of the finitary distribution monad, $\mathbf{K l}(D)$, which we will need in order to model Bayesian games (discussed in Section 4). Kl( $D$ ) inherits a copy/delete comonoid from Set, but its comonoid homomorphisms are the deterministic maps (i.e. precisely the non-probabilistic maps).

In the next section we introduce coends, a piece of categorical machinery that allows for an elegant generalisation of concrete lenses to arbitrary symmetric monoidal categories. We call these generalised lenses coend lenses or, simply, lenses. We will first introduce the technical notion before magicking it away with a diagrammatic calculus that represents what is 'really' going on.

### 3.2 Co-wedges and Coends

Co-wedges are a variant of co-cones of natural transformations applying to functors that act both covariantly and contravariantly on an argument. In Section 2.1 we noted that lenses have both covariant and contravariant components. We will see that this behaviour can be described by coends, which are initial co-wedges. For extra motivation, discussion, and examples, we refer the reader to [27].

Definition 3.2.1 (Co-wedge). Let $F: \mathcal{C}^{\mathrm{op}} \times \mathcal{C} \rightarrow \mathcal{D}$ be a functor. A co-wedge $c: F \rightarrow \alpha$ is an object $\alpha: \mathcal{D}$ together with maps $\left\{c_{a}: F(a, a) \rightarrow \alpha \mid a: \mathcal{C}\right\}$ such that, for any morphism $f: a^{\prime} \rightarrow a$, the diagram

commutes.

Definition 3.2.2 (Coend). A coend is a couniversal co-wedge. Diagrammatically, the coend of a functor $F: \mathcal{C}^{\mathrm{op}} \times \mathcal{C} \rightarrow D$ is a co-wedge $\left\{c_{a}: F(a, a) \rightarrow \operatorname{coend}(F) \mid a: \mathcal{C}\right\}$ such that for any other co-wedge $\left\{d_{a}: F(a, a) \rightarrow \alpha \mid a: \mathcal{C}\right\}$ and morphism $f: a^{\prime} \rightarrow a$ the diagram

commutes for a unique morphism $h: \operatorname{coend}(F) \rightarrow \alpha$.
We adopt the integral notation for coends, writing

$$
\int^{a: \mathcal{C}} F(a, a)
$$

for coend $(F)$. We will make use of the fact that coends can be characterised by the following coequaliser.

Lemma 3.2.3. Let $F: C^{\mathrm{op}} \times \mathcal{C} \rightarrow \mathcal{D}$. If $\mathcal{D}$ is cocomplete and $\mathcal{C}$ is small, the coend $\int^{a: \mathcal{C}} F(a, a)$ is given by the coequaliser of the pair of arrows

$$
\coprod_{\substack{a, a^{\prime}: \mathcal{C} \\ f: a^{\prime} \rightarrow a}} F\left(a, a^{\prime}\right) \frac{F_{1}}{F_{2}} \coprod_{a: \mathcal{C}} F(a, a)
$$

where the $f: a^{\prime} \rightarrow a$ components of $F_{1}$ and $F_{2}$ are $F\left(f, a^{\prime}\right)$ and $F(a, f)$ respectively.
When $\mathcal{C}$ is not small (as it usually is not), we need to show directly that coends exist.

### 3.3 Coend lenses

Much of the material in this section is worked out in much greater detail in [36], which serves as a good standard reference for coend lenses. We first give an abstract definition of coend lenses, then provide some justification.

Definition 3.3.1 (Coend lens). Let $X, S, Y$, and $R$ be objects in a symmetric monoidal category $\mathcal{C}$. A coend lens $l:(X, S) \rightarrow(Y, R)$ is an element of the set

$$
\int^{A: \mathcal{C}} \mathcal{C}(X, A \otimes Y) \times \mathcal{C}(A \otimes R, S)
$$

We think of the coend in the above definition as acting as a kind of existential quantifier over the type variable $A$, followed by a quotient (to be described) over the resulting structure. That is, a coend lens $l:(X, S) \rightarrow(Y, R)$ consists of an equivalence relation over triples comprised of a choice of type $A$, a morphism $v: X \rightarrow A \otimes Y$, and another morphism $u: A \otimes R \rightarrow S$.

By Lemma 3.2.3 we can characterise coend lenses $(X, S) \rightarrow(Y, R)$ as the elements of a particular coequaliser. Moreover, coequalisers in Set are given by quotients. Unpacking the coequaliser explicitly, coend lenses $(X, S) \rightarrow(Y, R)$ are given by the set of triples of the form described above, quotiented by the equivalence relation generated by (i.e. the smallest equivalence relation containing)

$$
\left(\left(f \otimes \operatorname{id}_{Y}\right) \circ v, u\right) \sim\left(v, u \circ\left(f \otimes \operatorname{id}_{R}\right)\right)
$$

for all $A, B: \mathcal{C}$ and $f: A \rightarrow B$. In diagrammatic form, the pair

is related to the pair


We refer to the types $A$ and $B$ as bound types ( $B$ is bound in the first diagram, $A$ in the second). In Section 4, we will see that this bound type keeps track of correlations between random variables in the Kleisli category of the distribution monad.

In vague terms, two pairs of morphisms are related if one can get from one to the other by 'sliding' a morphism off the bound type of one morphism on to the bound type of the other. Given a pair of morphisms $(v: X \rightarrow A \otimes Y, u: A \otimes R \rightarrow S)$, we write $[v, u]$ for their equivalence class. When we need to talk explicitly about the bound type of $[v, u]$ we write $[A, v, u]$ to specify that the pair $(v, u)$ has bound type $A$. We also adopt the convention that $l=\left[A_{l}, l_{v}, l_{u}\right]$ where, as with concrete lenses, we say that $l_{v}$ is the view morphism and $l_{u}$ is the update morphism. We follow [36], taking the hint from the diagrammatic representation of the equivalence relation by representing a coend lens $[v, u]:(X, S) \rightarrow(Y, R)$ as


We usually omit the bound type in diagrams for clarity. The equivalence relation is then simply


The equivalence relation permits the cancelling of isomorphisms:


Many proofs in this section proceed by allowing symmetric monoidal structure to interact with coend structure as, for example, in the following diagram.


The formal foundations of this class of diagrams are investigated in [37].
Example 3.3.2 (Identity lens). The identity lens $\operatorname{id}_{(X, S)}:(X, S) \rightarrow(X, S)$ is given by $\left[I, \mathrm{id}_{X}\right.$ : $\left.X \rightarrow X, \operatorname{id}_{S}: S \rightarrow S\right]$. Diagrammatically,


Example 3.3.3. A pair of morphisms $(f: X \rightarrow Y, g: R \rightarrow S)$ is encoded by the coend lens $[I, f, g]:(X, S) \rightarrow(Y, R)$ :


Definition 3.3.4 (Sequential composition of coend lenses). Let $[v, u]:(X, S) \rightarrow(Y, R)$ and $\left[v^{\prime}, u^{\prime}\right]:(Y, R) \rightarrow(Z, Q)$ be coend lenses. Define $\left[v^{\prime}, u^{\prime}\right] \circ[v, u]:(X, S) \rightarrow(Z, Q)$ to be


Explicitly,

$$
\left[A^{\prime}, v^{\prime}, u^{\prime}\right] \circ[A, v, u]=\left[A \otimes A^{\prime},\left(v^{\prime} \otimes \operatorname{id}_{A}\right) \circ v, u \circ\left(\operatorname{id}_{A} \otimes u^{\prime}\right)\right] .
$$

Theorem 3.3.5 (Coend lenses form a category). Suppose $\mathcal{C}$ is a monoidal category such that, for all objects $X, S, Y, R \in \mathcal{C}$,

$$
\int^{A: \mathcal{C}} \mathcal{C}(X, A \otimes Y) \times \mathcal{C}(A \otimes R, S)
$$

exists. Then there is a category $\mathbf{L e n s}_{\mathcal{C}}$ whose objects are pairs of objects in $\mathcal{C}$ and where

$$
\operatorname{Lens}_{\mathcal{C}}((X, S),(Y, R))=\int^{A: \mathcal{C}} \mathcal{C}(X, A \otimes Y) \times \mathcal{C}(A \otimes R, S)
$$

When $\mathcal{C}$ is small, the existence of sets of coend lenses of each type is guaranteed by the cocompleteness of Set. When $\mathcal{C}$ is not small, and the lens types correspond to coends indexed by a large category, we must verify that these sets exist by some other means (by, for example, giving a Set isomorphism). Fortunately, this is not difficult for the categories of interest in this work.

Definition 3.3.6 (Tensor composition of coend lenses). Let $[v, u]:(X, S) \rightarrow(Y, R)$ and $\left[v^{\prime}, u^{\prime}\right]$ : $\left(X^{\prime}, S^{\prime}\right) \rightarrow\left(Y^{\prime}, R^{\prime}\right)$ be coend lenses. Define $[v, u] \otimes\left[v^{\prime}, u^{\prime}\right]:\left(X \otimes X^{\prime}, S \otimes S^{\prime}\right) \rightarrow\left(Y \otimes Y^{\prime}, R \otimes R^{\prime}\right)$ to be


Explicitly, $[A, v, u] \otimes\left[A^{\prime}, v^{\prime}, u^{\prime}\right]$ is given by

$$
\left[\left(A \otimes A^{\prime}, \operatorname{id}_{A} \otimes s_{Y, A^{\prime}} \otimes \operatorname{id}_{Y^{\prime}}\right) \circ\left(v \otimes v^{\prime}\right),\left(u \otimes u^{\prime}\right) \circ\left(\operatorname{id}_{A} \otimes s_{A^{\prime}, R} \otimes \operatorname{id}_{R^{\prime}}\right)\right] .
$$

Theorem 3.3.7 $\left(\operatorname{Lens}_{\mathcal{C}}\right.$ is symmetric monoidal). The category $\operatorname{Lens}_{\mathcal{C}}$ is symmetric monoidal with the tensor given in Definition 3.3.6, monoidal unit $I=\left(I_{\mathcal{C}}, I_{\mathcal{C}}\right)$, and with structural morphisms inherited from $\mathcal{C}$ given by

$$
\begin{aligned}
\alpha_{(X, A),(Y, B),(Z, C)} & =\left[\alpha_{X, Y, Z}, \alpha_{A, B, C}^{-1}\right] \\
\lambda_{(X, A)} & =\left[\lambda_{X}, \lambda_{A}^{-1}\right] \\
\rho_{(X, A)} & =\left[\rho_{X}, \rho_{A}^{-1}\right] \\
s_{(X, A),(Y, B)} & =\left[s_{X, Y}, s_{B, A}\right] .
\end{aligned}
$$

Lemma 3.3.8. Lensset is isomorphic to CL. (More generally, when $\otimes$ is cartesian, Lens ${ }_{C}$ is isomorphic to an appropriately generalised definition of $\mathbf{C L}_{\mathcal{C}}$.)

### 3.4 Towards generalising open games

We could, at this point, attempt to define a (generalised) open game $\mathcal{G}:(X, S) \rightarrow(Y, R)$ over a symmetric monoidal category $\mathcal{C}$ as

1. A set $\Sigma$ of strategies;
2. A play function P : $\Sigma \rightarrow \operatorname{Lens}_{\mathcal{C}}((X, S),(Y, R))$; and
3. A best-response function $\mathrm{B}: \mathcal{C}(I, X) \times \mathcal{C}(Y, R) \rightarrow \operatorname{Rel}(\Sigma)$.

Call such generalised open games interim open games (for they will not live long). Sequential composition and tensor composition of interim open games could be defined much as we did for concrete open games. The problems begin to arise when one attempts to prove that this definition results in a symmetric monoidal category.

In proving that the associator was natural in $\mathbf{C L}$, we used the fact that the monoidal tensor in Set is cartesian. If the tensor of $\mathcal{C}$ is not cartesian, the local context of $\mathcal{G}$ in $\mathcal{G} \otimes(\mathcal{H} \otimes \mathcal{K})$ is different to the local context of $\mathcal{G}$ in $(\mathcal{G} \otimes \mathcal{H}) \otimes \mathcal{K}$. Let

$$
\left.\left.\begin{array}{rl}
\mathcal{G} & :\left(X_{1}, S_{1}\right) \\
\mathcal{H} & \rightarrow\left(Y_{1}, R_{1}\right) \\
\mathcal{K} & :\left(X_{3}, S_{3}\right)
\end{array}\right) \rightarrow\left(Y_{2}, R_{2}\right), R_{3}\right)
$$

be interim open games, $p \in \mathcal{C}\left(I, X_{1} \otimes X_{2} \otimes X_{3}\right)$, $k \in \mathcal{C}\left(Y_{1} \otimes Y_{2} \otimes Y_{3}, R_{1} \otimes R_{2} \otimes R_{3}\right)$, and $(\sigma, \tau, \mu) \in$ $\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}} \times \Sigma_{\mathcal{K}}$. The local context of $\mathcal{G}$ in $\mathcal{G} \otimes(\mathcal{H} \otimes \mathcal{K})$ is given by

whilst the local context of $\mathcal{G}$ in $(\mathcal{G} \otimes \mathcal{H}) \otimes \mathcal{K}$ is given by


In general, these morphisms are not the same. In the case where $\mathcal{C}$ is the Kleisli category of the distribution monad, the first morphism contains information about correlations between the types $X_{2}$ and $X_{3}$ whilst the second morphism does not. Consequently, the distinction between these two local contexts for $\mathcal{G}$ is substantive. Fortunately, coend lenses also provide a solution to this problem.

The high-level approach for defining a category of generalised open games is to use as few 'deleting' maps as possible. We do this by 'hiding' information in the bound variable of a coend lens whenever we would otherwise delete it. A consequence of this approach is that the correct definition of a 'context' for generalised open games is quite abstract, but we will see that this abstractness allows for more elegant proofs and, in any case, disappears when dealing with the categories we are actually interested in.

### 3.5 States, continuations, and contexts

In this section we define a generalised notion of context for open games. Observe that a state $\left[s, s^{\prime}\right] \in \operatorname{Lens}_{\mathcal{C}}(I,(X, S))$ has the form


More verbosely, a state $s \in \operatorname{Lens}_{\mathcal{C}}(I,(X, S))$ is the equivalence class of a choice of type $A: \mathcal{C}$ together with a state $s: \mathcal{C}(I, A \otimes X)$ in $\mathcal{C}$ and an effect $s^{\prime}: \mathcal{C}(A \otimes S, I)$ in $\mathcal{C}$. A useful interpretation of states in Lens $\mathcal{C}_{\mathcal{C}}$ is as a history/cohistory pair. (Cohistories are not yet well understood. They make proofs easier, but vanish in categories which make game-theoretic sense.)

An effect $\left[e, e^{\prime}\right] \in \operatorname{Lens}_{\mathcal{C}}((Y, R), I)$ has the form


Concerning effects, we have the following result.
Lemma 3.5.1. $\mathcal{C}(Y, R) \cong \operatorname{Lens}_{\mathcal{C}}((Y, R), I)$
Proof. The isomorphism $i: \mathcal{C}(Y, R) \rightarrow \operatorname{Lens}_{\mathcal{C}}((Y, R), I)$ is given by

$$
i(f: Y \rightarrow R)=\left[R, f, \operatorname{id}_{R}\right]=\left[Y, \mathrm{id}_{Y}, f\right]
$$

This result captures the idea that 'effects in Lens ${ }_{\mathcal{C}}$ are outcome functions in $\mathcal{C}$ '.
We can now define (generalised) contexts which consist of a coend over a state in Lens ${ }_{C}$ (a history/cohistory pair) and an effect in Lens $\mathcal{C}_{\mathcal{C}}$ (an outcome function). Contexts are therefore members of a double coend. This double coend turns out to be a state in the double lens category Lens $_{\text {Lens. }_{C}}$. From a purely technical standpoint, using double lenses allows for elegant proofs. From a heuristic perspective, we will see that the extra bound variable the double lens affords us enables us, in the case $\mathcal{C}=\mathbf{K l}(D)$, to store information about correlations between variables where we would otherwise have to take marginals.

Definition 3.5.2 (Context functor). The context functor $\mathbb{C}$ : Lens $_{\mathcal{C}} \times$ Lens $_{\mathcal{C}}{ }^{\mathrm{op}} \rightarrow$ Set is given by

$$
\begin{aligned}
\mathbb{C}(\Phi, \Psi) & =\int^{\Theta: \text { Lens }_{\mathcal{C}}} \operatorname{Lens}_{\mathcal{C}}(I, \Theta \otimes \Phi) \times \operatorname{Lens}_{\mathcal{C}}(\Theta \otimes \Psi, I) \\
& =\operatorname{Lens}_{\mathbf{L e n s}_{\mathcal{C}}}(I,(\Phi, \Psi))
\end{aligned}
$$

Elements of $\mathbb{C}(\Phi, \Psi)$ are called contexts.
(We use the letters $\Phi, \Psi$ to refer to objects of $\operatorname{Lens}_{\mathcal{C}}$, which are pairs of objects of $\mathcal{C}$.)
As a context $[p, k] \in \mathbb{C}(\Phi, \Psi)$ is just a state in Lens $_{\text {Lens }_{\mathcal{C}}}$, so it admits a graphical representation as


This is neat, and means many of the results in the rest of this section can be carried out graphically.

### 3.6 General open games

We have now arrived at a level of generality where we can define generalised open games in a way that is obviously analogous to concrete open games. Given $\Phi, \Psi \in \operatorname{Lens}_{\mathcal{C}}$, an open game consists of a set of strategy profiles, a family of lenses indexed by the set of strategy profiles, and a best-response function which takes a context as input and returns a relation on strategy profiles.

Definition 3.6.1 (Open game). Let $\Phi, \Psi \in$ Lens $_{\mathcal{C}}$. An open game $\mathcal{G}: \Phi \rightarrow \Psi$ consists of

1. A set of strategy profiles $\Sigma$;
2. A play function $\mathrm{P}: \Sigma \rightarrow \operatorname{Lens}_{\mathcal{C}}(\Phi, \Psi)$; and
3. A best-response function $\mathrm{B}: \mathbb{C}(\Phi, \Psi) \rightarrow \operatorname{Rel}(\Sigma)$.

The rationale here is much the same as it is with concrete open games. The play function takes a strategy profile as input and returns a lens describing an open play of the game. Best response takes a context as argument that provides the information necessary for the game to make informed strategic decisions, and returns a relation on strategies.

As with concrete open games, we define a notion of atomic open game:
Definition 3.6.2. An atomic open game $a: \Phi \rightarrow \Psi$ is an open game such that

1. $\Sigma_{a} \subseteq \operatorname{Lens}_{\mathcal{C}}(\Phi, \Psi)$;
2. For all $l \in \Sigma_{a}, \mathrm{P}_{a}(l)=l$; and
3. For all contexts $c \in \mathbb{C}(\Phi, \Psi), \mathrm{B}_{a}(c)$ is constant.

Atomic open games are uniquely specified by a subset $\Sigma \subseteq \operatorname{Lens}_{\mathcal{C}}(\Phi, \Psi)$ and a selection function B: $\mathbb{C}(\Phi, \Psi) \rightarrow \mathcal{P}(\Sigma)$, and we will sometimes specify atomic open games via this data. We refer to atomic open games simply as atoms.

Example 3.6.3. The identity atom $\operatorname{id}_{\Phi}: \Phi \rightarrow \Phi$ is given by $\Sigma=\left\{\operatorname{id}_{\Phi}\right\}, \mathrm{B}(c)=\left\{\operatorname{id}_{\Phi}\right\}$ for all $c \in \mathbb{C}(\Phi, \Phi)$.

Example 3.6.4 (Computation). Let $f: \mathcal{C}(X, Y)$ and $g: \mathcal{C}(R, S)$ be morphisms in $\mathcal{C}$. Define the atom $\langle f, g\rangle:(X, S) \rightarrow(Y, R)$ by

1. $\Sigma_{\langle f, g\rangle}=\{[f, g]\}$; and
2. $\mathrm{B}_{\langle f, g\rangle}(c)=\{[f, g]\}$ for all $c \in \mathbb{C}((X, S),(Y, R))$.

### 3.7 Composing open games

The heuristic for sequential composition of general open games is much the same as for concrete open games in Subsection 2.7. The only difference is that we are now using coend lenses rather than concrete lenses, and contexts also are slightly different. Best response of a sequential composite $\mathcal{H} \circ \mathcal{G}$ is still defined by forming local contexts for $\mathcal{G}$ and $\mathcal{H}$.

Definition 3.7.1 (Sequential composition of open games). Let $\mathcal{G}: \Phi \rightarrow \Psi$ and $\mathcal{H}: \Psi \rightarrow \Xi$ be open games. Define $\mathcal{H} \circ \mathcal{G}: \Phi \rightarrow \Xi$ by

1. $\Sigma_{\mathcal{H} \circ \mathcal{G}}=\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}$,
2. $\mathrm{P}_{\mathcal{H} \circ \mathcal{G}}(\sigma, \tau)=\mathrm{P}_{\mathcal{H}}(\tau) \circ \mathrm{P}_{\mathcal{G}}(\sigma)$,
3. $\mathrm{B}_{\mathcal{H} \circ \mathcal{G}}([p, k])(\sigma, \tau)=\mathrm{B}_{\mathcal{G}}\left(\left[p, k \circ \mathrm{P}_{\mathcal{H}}(\tau)\right]\right)(\sigma) \times \mathrm{B}_{\mathcal{H}}\left(\left[\mathrm{P}_{\mathcal{G}}(\sigma) \circ p, k\right]\right)(\tau)$.

Given a context $[p, k] \in \mathbb{C}(\Phi, \Xi)$ represented by the diagram

the local context for $\mathcal{G}$ given a strategy $\tau \in \Sigma_{\mathcal{H}}$ is given by

and given a strategy $\sigma \in \Sigma_{\mathcal{G}}$ the local context for $\mathcal{H}$ is given by


In this representation the process of taking a local context is non-arbitrary, and obviously associative.

### 3.8 The tensor of open games

Again, the heuristic for defining the tensor of open games is much as it was for concrete open games. We will first formalise the notion of 'local context' for tensored general open games.

Definition 3.8.1 (Local contexts for tensor composition). Define the left local context function

$$
\mathcal{L}_{\Phi, \Phi^{\prime}, \Psi, \Psi^{\prime}}: \mathbb{C}\left(\Phi \otimes \Phi^{\prime}, \Psi \otimes \Psi^{\prime}\right) \times \operatorname{Lens}_{\mathcal{C}}\left(\Phi^{\prime}, \Psi^{\prime}\right) \rightarrow \mathbb{C}(\Phi, \Psi)
$$

by


Define the right local context function

$$
\mathcal{R}_{\Phi, \Phi^{\prime}, \Psi, \Psi^{\prime}}: \mathbb{C}\left(\Phi \otimes \Phi^{\prime}, \Psi \otimes \Psi^{\prime}\right) \times \operatorname{Lens}_{\mathcal{C}}(\Phi, \Psi) \rightarrow \mathbb{C}\left(\Phi^{\prime}, \Psi^{\prime}\right)
$$

by


We will usually suppress the subscripts of $\mathcal{L}$ and $\mathcal{R}$ as the types can be inferred from context.
Definition 3.8.2 (Tensor composition of open games). Let $\mathcal{G}: \Phi \rightarrow \Psi$ and $\mathcal{H}: \Phi^{\prime} \rightarrow \Psi^{\prime}$ be open games. Define $\mathcal{G} \otimes \mathcal{H}: \Phi \otimes \Phi^{\prime} \rightarrow \Psi \otimes \Psi^{\prime}$ by

- $\Sigma_{\mathcal{G} \otimes \mathcal{H}}=\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}} ;$
- $\mathrm{P}_{\mathcal{G} \otimes \mathcal{H}}(\sigma, \tau)=\mathrm{P}_{\mathcal{G}}(\sigma) \otimes \mathrm{P}_{\mathcal{H}}(\tau)\left(\right.$ in Lens $\left._{\mathcal{C}}\right) ;$
- Define $\mathrm{B}_{\mathcal{G} \otimes \mathcal{H}}: \mathbb{C}\left(\Phi \otimes \Phi^{\prime}, \Psi \otimes \Psi^{\prime}\right) \rightarrow \operatorname{Rel}\left(\Sigma_{\mathcal{G} \otimes \mathcal{H}}\right)$ by

$$
\mathrm{B}_{\mathcal{G} \otimes \mathcal{H}}(c)(\sigma, \tau)=\mathrm{B}_{\mathcal{G}}\left(\mathcal{L}\left(c, \mathrm{P}_{\mathcal{H}}(\tau)\right)\right)(\sigma) \times \mathrm{B}_{\mathcal{H}}\left(\mathcal{R}\left(c, \mathrm{P}_{\mathcal{G}}(\sigma)\right)\right)(\tau)
$$

### 3.9 Equivalence of open games

As in Section 2.9, we need to quotient open games in order to obtain a category.
Definition 3.9.1 (Isomorphism of open games). Let $\mathcal{G}, \mathcal{H}: \Phi \rightarrow \Psi$ be open games. An isomorphism of open games $\alpha: \mathcal{G} \rightarrow \mathcal{H}$ is a bijection $\alpha: \Sigma_{\mathcal{G}} \rightarrow \Sigma_{\mathcal{H}}$ such that

1. $\mathrm{P}_{\mathcal{G}}(\sigma)=\mathrm{P}_{\mathcal{H}}(\alpha(\sigma))$ for all $\sigma \in \Sigma_{\mathcal{G}}$; and
2. For all $\sigma, \sigma^{\prime} \in \Sigma_{\mathcal{G}}$ and $c \in \mathbb{C}(\Phi, \Psi),\left(\sigma, \sigma^{\prime}\right) \in \mathrm{B}_{\mathcal{G}}(c)$ iff $\left(\alpha(\sigma), \alpha\left(\sigma^{\prime}\right)\right) \in \mathrm{B}_{\mathcal{H}}(c)$.

Definition 3.9.2 (Equivalence of open games). Let $\mathcal{G}, \mathcal{H}: \Phi \rightarrow \Psi$ be open games. $\mathcal{G}$ and $\mathcal{H}$ are equivalent, written $\mathcal{G} \sim \mathcal{H}$, if there exists an isomorphism $\alpha: \mathcal{G} \rightarrow \mathcal{H}$. We write $[\mathcal{G}]$ for the equivalence class of $\mathcal{G}$ under this relation.

Lemma 3.9.3. Let $\mathcal{G}, \mathcal{G}^{\prime}: \Phi \rightarrow \Psi, \mathcal{H}, \mathcal{H}^{\prime}: \Psi \rightarrow \Xi$, and $\mathcal{K}, \mathcal{K}^{\prime}: \Phi^{\prime} \rightarrow \Psi^{\prime}$ be open games. Then

1. If $\mathcal{G} \sim \mathcal{G}^{\prime}$ and $\mathcal{H} \sim \mathcal{H}^{\prime}$, then $\mathcal{H} \circ \mathcal{G} \sim \mathcal{H}^{\prime} \circ \mathcal{G}^{\prime}$; and
2. If $\mathcal{G} \sim \mathcal{G}^{\prime}$ and $\mathcal{K} \sim \mathcal{K}^{\prime}$, then $\mathcal{G} \otimes \mathcal{K} \sim \mathcal{G}^{\prime} \otimes \mathcal{K}^{\prime}$.

Demonstrating equivalence in the cases of interest will always be trivial, and so we simply specify the witnessing bijection between strategy sets.

### 3.10 The category of open games

That equivalence classes of open games form a category follows easily from the fact that coend lenses form a category.

Lemma 3.10.1. Sequential composition of equivalence classes of open games is associative.
Proof. Suppose we have open games

$$
\Phi \xrightarrow{\mathcal{G}} \Psi \xrightarrow{\mathcal{H}} \Xi \xrightarrow{\mathcal{K}} \Upsilon .
$$

The equivalence between $(\mathcal{K} \circ \mathcal{H}) \circ \mathcal{G}$ and $\mathcal{K} \circ(\mathcal{H} \circ \mathcal{G})$ will be witnessed by the isomorphism $\beta: \Sigma_{\mathcal{G}} \times\left(\Sigma_{\mathcal{H}} \times \Sigma_{\mathcal{K}}\right) \rightarrow\left(\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}\right) \times \Sigma_{\mathcal{K}},(\sigma,(\tau, \mu)) \mapsto((\sigma, \tau), \mu)$. Let $\sigma \in \Sigma_{\mathcal{G}}, \tau \in \Sigma_{\mathcal{H}}$, and $\mu \in \Sigma_{\mathcal{K}}$. Then $\mathrm{P}_{(\mathcal{K} \circ \mathcal{H}) \circ \mathcal{G}}(\sigma,(\tau, \mu))=\mathrm{P}_{\mathcal{K} \circ(\mathcal{H} \circ \mathcal{G})}((\sigma, \tau), \mu)$ by associativity of composition in Lens ${ }_{\mathcal{C}}$. Let $[p, k] \in \mathbb{C}(\Phi, \Upsilon)$ be a context. Then

$$
\begin{aligned}
& \left((\sigma,(\tau, \mu)),\left(\sigma^{\prime},\left(\tau^{\prime}, \mu^{\prime}\right)\right)\right) \in \mathrm{B}_{\left(\mathcal{K}^{\prime \mathcal{H})} \circ \mathcal{G}\right.}([p, k]) \\
\Longleftrightarrow & \left(\sigma, \sigma^{\prime}\right) \in \mathrm{B}_{\mathcal{G}}\left(\left[p, k \circ \mathrm{P}_{\mathcal{K}}(\mu) \circ \mathrm{P}_{\mathcal{H}}(\tau)\right]\right) \text { and }\left(\tau, \tau^{\prime}\right) \in \mathrm{B}_{\mathcal{H}}\left(\left[\mathrm{P}_{\mathcal{G}}(\sigma) \circ p, k \circ \mathrm{P}_{\mathcal{K}}(\mu)\right]\right) \\
& \quad \text { and }\left(\mu, \mu^{\prime}\right) \in \mathrm{B}_{\mathcal{K}}\left(\left[\mathrm{P}_{\mathcal{H}}(\tau) \circ \mathrm{P}_{\mathcal{G}}(\sigma) \circ p, k\right]\right) \\
\Longleftrightarrow & \left(((\sigma, \tau), \mu),\left(\left(\sigma^{\prime}, \tau^{\prime}\right), \mu^{\prime}\right)\right) \in \mathrm{B}_{\mathcal{K}_{\circ}(\mathcal{H} \circ \mathcal{G})}([p, k])
\end{aligned}
$$

Theorem 3.10.2. If $\mathbf{L e n s}_{\mathcal{C}}$ exists, there exists a category $\mathbf{G a m e}_{\mathcal{C}}$ with pairs of objects in $\mathcal{C}$ as objects and equivalence classes of open games as morphisms.

Proof. All that remains to be checked is that the identity computation defined in Example 3.6.3 is an identity morphism, and this follows from easy checks.

### 3.11 The symmetric monoidal structure of open games

We now prove that $\otimes$ is functorial. The proof is a good demonstration of the utility of coend diagrams. In the commutative squares in the following lemma, the top path describes how local contexts are formed in, say, $\left(\mathcal{H} \otimes \mathcal{H}^{\prime}\right) \circ\left(\mathcal{G} \otimes \mathcal{G}^{\prime}\right)$ and the bottom path describes how local contexts are formed in $(\mathcal{H} \circ \mathcal{G}) \otimes\left(\mathcal{H}^{\prime} \circ \mathcal{G}^{\prime}\right)$. That the squares commute follows by inspection of the appropriate coend diagrams.

Lemma 3.11.1. Suppose we have coend lenses

$$
\begin{aligned}
& \Phi \xrightarrow{l} \Psi \xrightarrow{m} \Xi \\
& \Phi^{\prime} \xrightarrow{l^{\prime}} \Psi^{\prime} \xrightarrow{m^{\prime}} \Xi^{\prime}
\end{aligned}
$$

The following diagrams commute:
1.

2.

3.

4.


Proof. The four squares are given respectively by the following equalities of coend diagrams:
1.


2.

$=$

3.

$=$

4.



Functoriality of the tensor in $\mathbf{G a m e}_{\mathcal{C}}$ then follows easily.
Corollary 3.11.2. $\otimes:$ Game $_{\mathcal{C}} \times$ Game $_{\mathcal{C}} \rightarrow$ Game $_{\mathcal{C}}$ is a functor.
Proof. Suppose we have open games


Note that $\Sigma_{(\mathcal{H} \circ \mathcal{G}) \otimes\left(\mathcal{H}^{\prime} \circ \mathcal{G}^{\prime}\right)}=\left(\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}\right) \times\left(\Sigma_{\mathcal{G}^{\prime}} \times \Sigma_{\mathcal{H}^{\prime}}\right)$ and $\Sigma_{\left(\mathcal{H} \otimes \mathcal{H}^{\prime}\right) \circ\left(\mathcal{G} \otimes \mathcal{G}^{\prime}\right)}=\left(\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{G}^{\prime}}\right) \times\left(\Sigma_{\mathcal{H}} \times \Sigma_{\mathcal{H}^{\prime}}\right)$. The isomorphism $\beta:\left(\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}\right) \times\left(\Sigma_{\mathcal{G}^{\prime}} \times \Sigma_{\mathcal{H}^{\prime}}\right) \rightarrow\left(\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{G}^{\prime}}\right) \times\left(\Sigma_{\mathcal{H}} \times \Sigma_{\mathcal{H}^{\prime}}\right)$ witnessing the equivalence between $(\mathcal{H} \circ \mathcal{G}) \otimes\left(\mathcal{H}^{\prime} \circ \mathcal{G}^{\prime}\right)$ and $\left(\mathcal{H} \otimes \mathcal{H}^{\prime}\right) \circ\left(\mathcal{G} \otimes \mathcal{G}^{\prime}\right)$ is given by $\left((\sigma, \tau),\left(\sigma^{\prime}, \tau^{\prime}\right)\right) \mapsto\left(\left(\sigma, \sigma^{\prime}\right),\left(\tau, \tau^{\prime}\right)\right)$. Lens $_{\mathcal{C}}$ is symmetric monoidal and, hence,

$$
\mathrm{P}_{(\mathcal{H} \circ \mathcal{G}) \otimes\left(\mathcal{H}^{\prime} \otimes \mathcal{G}^{\prime}\right)}\left((\sigma, \tau),\left(\sigma^{\prime}, \tau^{\prime}\right)\right)=\mathrm{P}_{\left(\mathcal{H} \otimes \mathcal{H}^{\prime}\right) \circ(\mathcal{G} \otimes \mathcal{G})}\left(\left(\sigma, \sigma^{\prime}\right),\left(\tau, \tau^{\prime}\right)\right) .
$$

Using Lemma 3.11.1,

$$
\begin{aligned}
& \left(\left((\sigma, \tau),\left(\sigma^{\prime}, \tau^{\prime}\right)\right),\left(\left(\sigma^{\prime \prime}, \tau^{\prime \prime}\right),\left(\sigma^{\prime \prime \prime}, \tau^{\prime \prime \prime}\right)\right)\right) \in \mathrm{B}_{(\mathcal{H} \circ \mathcal{G})} \otimes\left(\mathcal{H}^{\prime} \circ \mathcal{G}^{\prime}\right) \\
\Longleftrightarrow & \left(\sigma, \sigma^{\prime \prime}\right) \in \mathrm{B}_{\mathcal{G}}\left(\mathbb{C}\left(\Phi, \mathrm{P}_{\mathcal{H}}(\tau)\right) \circ \mathcal{L}\left(-, \mathrm{P}_{\mathcal{H}^{\prime}}\left(\tau^{\prime}\right) \circ \mathrm{P}_{\mathcal{G}^{\prime}}\left(\sigma^{\prime}\right)\right)(c)\right) \\
& \text { and }\left(\tau, \tau^{\prime \prime}\right) \in \mathrm{B}_{\mathcal{H}}\left(\mathbb{C}\left(\mathrm{P}_{\mathcal{G}}(\sigma), \Xi\right) \circ \mathcal{L}\left(-, \mathrm{P}_{\mathcal{H}^{\prime}}\left(\tau^{\prime}\right) \circ \mathrm{P}_{\mathcal{G}^{\prime}}\left(\sigma^{\prime}\right)\right)(c)\right) \\
& \text { and }\left(\sigma^{\prime}, \sigma^{\prime \prime \prime}\right) \in \mathrm{B}_{\mathcal{G}^{\prime}}\left(\mathbb{C}\left(\Phi^{\prime}, \mathrm{P}_{\mathcal{H}^{\prime}}\left(\tau^{\prime}\right)\right) \circ \mathcal{R}\left(-, \mathrm{P}_{\mathcal{H}}(\tau) \circ \mathrm{P}_{\mathcal{H}}(\sigma)\right)(c)\right) \\
& \text { and }\left(\tau^{\prime}, \tau^{\prime \prime \prime}\right) \in \mathrm{B}_{\mathcal{H}^{\prime}}\left(\mathbb{C}\left(\mathrm{P}_{\mathcal{G}^{\prime}}\left(\sigma^{\prime}\right), \Xi^{\prime}\right) \circ \mathcal{R}\left(i, \mathrm{P}_{\mathcal{H}}(\tau) \circ \mathrm{P}_{\mathcal{H}}(\sigma)\right)(c)\right) \\
\Longleftrightarrow & \left(\sigma, \sigma^{\prime \prime}\right) \in \mathrm{B}_{\mathcal{G}}\left(\mathcal{L}\left(-, \mathrm{P}_{\mathcal{G}^{\prime}}\left(\sigma^{\prime}\right)\right) \circ \mathbb{C}\left(\Phi \otimes \Phi^{\prime}, \mathrm{P}_{\mathcal{H}}(\tau) \otimes \mathrm{P}_{\mathcal{H}^{\prime}}\left(\tau^{\prime}\right)\right)(c)\right) \\
& \text { and }\left(\sigma^{\prime}, \sigma^{\prime \prime \prime}\right) \in \mathrm{B}_{\mathcal{G}^{\prime}}\left(\mathcal{R}\left(-, \mathrm{P}_{\mathcal{G}}(\sigma)\right) \circ \mathbb{C}\left(\Phi \otimes \Phi^{\prime}, \mathrm{P}_{\mathcal{H}}(\tau) \otimes \mathrm{P}_{\mathcal{H}^{\prime}}\left(\tau^{\prime}\right)\right)(c)\right) \\
& \text { and }\left(\tau, \tau^{\prime \prime}\right) \in \mathrm{B}_{\mathcal{H}}\left(\mathcal{L}\left(-, \mathrm{P}_{\mathcal{H}^{\prime}}\left(\tau^{\prime}\right)\right) \circ \mathbb{C}\left(\mathrm{P}_{\mathcal{G}}(\sigma) \otimes \mathrm{P}_{\mathcal{G}^{\prime}}\left(\sigma^{\prime}\right), \Xi \otimes \Xi^{\prime}\right)(c)\right) \\
& \text { and }\left(\tau^{\prime}, \tau^{\prime \prime \prime}\right) \in \mathrm{B}_{\mathcal{H}^{\prime}}\left(\mathcal{R}\left(-, \mathrm{P}_{\mathcal{H}}(\tau)\right) \circ \mathbb{C}\left(\mathrm{P}_{\mathcal{G}}(\sigma) \otimes \mathrm{P}_{\mathcal{G}^{\prime}}\left(\sigma^{\prime}\right), \Xi \otimes \Xi^{\prime}\right)(c)\right) \\
\Longleftrightarrow & \left(\left(\left(\sigma, \sigma^{\prime}\right),\left(\tau, \tau^{\prime}\right)\right),\left(\left(\sigma^{\prime \prime}, \sigma^{\prime \prime \prime}\right),\left(\tau^{\prime \prime}, \tau^{\prime \prime \prime}\right)\right)\right) \in \mathrm{B}_{\left(\mathcal{H} \otimes \mathcal{H}^{\prime}\right) \circ\left(\mathcal{G} \otimes \mathcal{G}^{\prime}\right)}(c) .
\end{aligned}
$$

Definition 3.11.3. The structural isomorphisms in $\mathbf{G a m e}_{\mathcal{C}}$ are given by

$$
\begin{aligned}
\alpha_{(X, A),(Y, B),(Z, C)} & =\left\langle\alpha_{X, Y, Z}, \alpha_{A, B, C}^{-1}\right\rangle \\
\rho_{(X, A)} & =\left\langle\rho_{X}, \rho_{A}^{-1}\right\rangle \\
\lambda_{(X, A)} & =\left\langle\lambda_{X}, \lambda_{A}^{-1}\right\rangle \\
s_{(X, A),(Y, B)} & =\left\langle s_{X, Y}, s_{B, Y}\right\rangle .
\end{aligned}
$$

Lemma 3.11.4. The structural isomorphisms are natural in Game $_{\mathcal{C}}$.

Proof. We show that the associator is natural. Naturality of the other stuctural maps follow by similar arguments. Let $\mathcal{G}_{i}: \Phi_{i} \rightarrow \Psi_{i}$ for $i \in\{1,2,3\}$. Note that $\Sigma_{\alpha \circ\left(\mathcal{G}_{1} \otimes\left(\mathcal{G}_{2} \otimes \mathcal{G}_{3}\right)\right)}=\left(\Sigma_{\mathcal{G}_{1}} \times\right.$ $\left.\left(\Sigma_{\mathcal{G}_{2}} \times \Sigma_{\mathcal{G}_{3}}\right)\right) \times\{\alpha\}$ and $\Sigma_{\left(\left(\mathcal{G}_{1} \otimes \mathcal{G}_{2}\right) \otimes \mathcal{G}_{3}\right) \circ \alpha}=\{\alpha\} \times\left(\left(\Sigma_{\mathcal{G}_{1}} \times \Sigma_{\mathcal{G}_{2}}\right) \times \Sigma_{\mathcal{G}_{3}}\right)$. The equivalence between $\alpha \circ\left(\mathcal{G}_{1} \otimes\left(\mathcal{G}_{2} \otimes \mathcal{G}_{3}\right)\right)$ and $\left(\left(\mathcal{G}_{1} \otimes \mathcal{G}_{2}\right) \otimes \mathcal{G}_{3}\right) \circ \alpha$ will be witnessed by the isomorphism $((\sigma,(\tau, \mu)), \alpha) \mapsto$ $(\alpha,((\sigma, \tau), \mu))$. Let $\sigma \in \Sigma_{\mathcal{G}_{1}}, \tau \in \Sigma_{\mathcal{G}_{2}}, \mu \in \Sigma_{\mathcal{G}_{3}}$, and $[p, k] \in \mathbb{C}\left(\left(\Phi_{1} \otimes\left(\Phi_{2} \otimes \Phi_{3}\right)\right),\left(\left(\Psi_{1} \otimes \Psi_{2}\right) \otimes \Psi_{3}\right)\right)$. We note that the local context for $\mathcal{G}_{1}$ given this data is the same for both games. The local context of $\mathcal{G}_{1}$ is given by

in $\alpha \circ\left(\mathcal{G}_{1} \otimes\left(\mathcal{G}_{2} \otimes \mathcal{G}_{3}\right)\right)$ and by

in $\left(\left(\mathcal{G}_{1} \otimes \mathcal{G}_{2}\right) \otimes \mathcal{G}_{3}\right) \circ \alpha$. This two morphisms are evidently equal. Similar diagrams demonstrate that the local contexts for $\mathcal{G}_{2}$ and $\mathcal{G}_{3}$ are the same in both games also.

Theorem 3.11.5. Game $_{\mathcal{C}}$ is symmetric monoidal.
Proof. All that remains to be shown is that the Mac Lane pentagon and triangle axioms are satisfied, but this follows easily as the underlying category $\mathcal{C}$ is symmetric monoidal.

### 3.12 Nice categories of open games

In this section we show how the notion of 'cohistory' collapses when the monoidal unit $I$ of the underlying monoidal category $\mathcal{C}$ is terminal. With cohistories gone, we will see that Game ${ }_{\mathcal{C}}$ has a very natural game-theoretic interpretation.

Lemma 3.12.1 ([36]). If the monoidal unit of $\mathcal{C}$ is terminal, then $\operatorname{Lens}_{\mathcal{C}}(I,(X, S)) \cong \mathcal{C}(I, X)$.
The isomorphism $i: \mathcal{C}(I, X) \rightarrow \operatorname{Lens}_{\mathcal{C}}(I,(X, S))$ is explicitly given by $p \mapsto\left[p,!_{s}\right]$. In a diagram,


The following fact appears as [27, exercise 1.13]; thanks to Guillaume Boisseau and Amar Hadzihasanovic for the discussion at [39].

Lemma 3.12.2. If $F \dashv U: \mathcal{D} \rightarrow \mathcal{C}$ is any adjunction and $G: \mathcal{D}^{\mathrm{op}} \times \mathcal{C} \rightarrow$ Set any functor, then

$$
\int^{C \in \mathcal{C}} G(F(C), C) \cong \int^{D \in \mathcal{D}} G(D, U(D))
$$

Proof.

$$
\begin{aligned}
\int^{C \in \mathcal{C}} G(F(C), C) & \cong \iint^{C \in \mathcal{C}, D \in \mathcal{D}} G(D, C) \times \mathcal{D}(F(C), D) & & \text { (ninja Yoneda lemma) } \\
& \cong \iint^{D \in \mathcal{D}, C \in \mathcal{C}} G(D, C) \times \mathcal{C}(C, U(D)) & & \text { (adjunction, Fubini theorem) } \\
& \cong \int^{D \in \mathcal{D}} G(D, U(D)) & & \text { (ninja Yoneda lemma) }
\end{aligned}
$$

Lemma 3.12.3. If the monoidal unit of $\mathcal{C}$ is terminal, then

$$
\mathbb{C}((X, S),(Y, R)) \cong \operatorname{Lens}_{\mathcal{C}}((I, R),(X, Y))
$$

Proof. Let $F: \mathcal{C} \rightarrow$ Lens $_{\mathcal{C}}$ be the embedding $F(X)=(X, I)$. When the monoidal unit of $\mathcal{C}$ is terminal, this functor has a right adjoint $U: \operatorname{Lens}_{\mathcal{C}} \rightarrow \mathcal{C}$ that is given on objects by $U(X, S)=X$. On morphisms $U$ is defined by the universal maps

$$
\int^{A \in \mathcal{C}} \mathcal{C}(X, A \otimes Y) \times \mathcal{C}(A \otimes R, S) \rightarrow \mathcal{C}(X, Y)
$$

induced by the dinatural (in $A$ ) maps $\mathcal{C}(X, A \otimes Y) \times \mathcal{C}(A \otimes R, S) \rightarrow \mathcal{C}(X, Y)$, taking ( $v, u$ ) to $X \xrightarrow{v} A \otimes Y \xrightarrow{!_{A} \otimes Y} I \otimes Y \cong Y$. The adjunction $F \dashv U$ is given by the natural isomorphism

$$
\int^{A \in \mathcal{C}} \mathcal{C}(X, A \otimes Y) \times \mathcal{C}(A \otimes R, I) \cong \int^{A \in \mathcal{C}} \mathcal{C}(X, A \otimes Y) \times \mathcal{C}(A, I) \cong \mathcal{C}(X, I \otimes Y) \cong \mathcal{C}(X, Y)
$$

In the previous lemma we take $G:$ Lens $_{\mathcal{C}}^{\text {op }} \times \mathcal{C} \rightarrow$ Set to be

$$
G\left(\left(\Phi, \Phi^{\prime}\right), \Theta\right)=\mathcal{C}(I, \Theta \otimes X) \times \operatorname{Lens}_{\mathcal{C}}\left(\left(\Phi, \Phi^{\prime}\right) \otimes(Y, R), I\right)
$$

Then

$$
\mathbb{C}((X, S),(Y, R)) \cong \int^{\left(\Theta, \Theta^{\prime}\right) \in \operatorname{Lens}_{\mathcal{C}}} G\left(\left(\Theta, \Theta^{\prime}\right), \Theta\right) \cong \int^{\Theta \in \mathcal{C}} G((\Theta, I), \Theta) \cong \operatorname{Lens}_{\mathcal{C}}((I, R),(X, Y))
$$

with the three isomorphisms respectively using Lemmas 3.12.1, 3.12.2 and 3.5.1.
In the case where the monoidal unit of $\mathcal{C}$ is terminal, the type of best response for an open game $\mathcal{G}:(X, S) \rightarrow(Y, R)$ is equivalently

$$
\mathrm{B}_{\mathcal{G}}: \operatorname{Lens}_{\mathcal{C}}((I, R),(X, Y)) \rightarrow \operatorname{Rel}\left(\Sigma_{\mathcal{G}}\right)
$$

We have seen that expressing contexts as states in the double lens category is a good level of abstraction for categories of open games, allowing for elegant diagrammatic proofs. From a gametheoretic perspective, however, it will make more sense to express contexts as equivalence classes $[p, k, \Theta]: \operatorname{Lens}_{\mathcal{C}}((I, R),(X, Y))$. This is because a state $p: \mathcal{C}(I, \Theta \otimes X)$ is easily seen to correspond to a history for an open game and the function $k: \Theta \otimes Y \rightarrow R$ acts like an outcome function. In this way, we can specify a context for an open game in much the same way as we did for concrete open games in section 2 .

The coend diagram

of a context $[p, k] \in \operatorname{Lens}_{\mathcal{C}}((I, R),(X, Y))$ neatly illustrates that a context is a game state with a 'hole' in it. If we think of a game $\mathcal{G}:(X, S) \rightarrow(Y, R)$ as a player in a larger game, then $p$ corresponds to the things that have happened in the game before $\mathcal{G}$ gets to act; $k$ corresponds to what will happen in the game after $\mathcal{G}$ acts; and the gap in the diagram corresponds to the part of
the game where $\mathcal{G}$ gets to influence the outcome. Alternatively, a context is that which becomes a game once $\mathcal{G}$ has decided which strategy to play, whereby playing that strategy will fill in the gap in the context.

Given open games $\mathcal{G}:(X, S) \rightarrow(Y, R), \mathcal{H}:(Y, R) \rightarrow(Z, Q)$, strategies $\sigma \in \Sigma_{\mathcal{G}}, \tau \in \Sigma_{\mathcal{H}}$, and a context $[p, k] \in \operatorname{Lens}_{\mathcal{C}}((I, R),(X, Z))$, the local context for $\mathcal{G}$ in $\mathcal{H} \circ \mathcal{G}$ is given by

and the local context for $\mathcal{H}$ is given by


Given another open game $\mathcal{K}:\left(X^{\prime}, S^{\prime}\right) \rightarrow\left(Y^{\prime}, R^{\prime}\right)$, a context $[p, k] \in \operatorname{Lens}_{\mathcal{C}}\left(\left(I, R \otimes R^{\prime}\right),(X \otimes\right.$ $\left.X^{\prime}, Y \otimes Y^{\prime}\right)$ ), and a strategy $\mu \in \Sigma_{\mathcal{H}}$, the local contexts for $\mathcal{G}$ and $\mathcal{H}$ in $\mathcal{G} \otimes \mathcal{H}$ are given by

and

respectively.

## 4 Bayesian open games

In this section we will zero in on open games with a specific lens structure. As we will show this class of open games will address the shortcomings of concrete open games.

### 4.1 Commutative monads

Recall that a monad $T$ over a monoidal category $\mathcal{C}$ is strong if it comes with a strength natural transformation $t_{A, B}: A \otimes T B \rightarrow T(A \otimes B)$ satisfying various coherence conditions.

We have the following result guaranteeing the existence of a large class of coend lens categories. We refer the reader to [36] for a much more in-depth discussion of the following result, and many more examples of when lens categories exist.

Theorem 4.1.1 ([36]). If $T$ is a strong monad, then $\mathbf{L e n s}_{\mathbf{K l}(T)}$ exists ${ }^{2}$.

[^2]Definition 4.1.2 (Commutative monad). Let $T$ be a strong monad with strength $t$ over a monoidal category $\mathcal{C}$. Define the costrength natural transformation $t_{A, B}^{\prime}: T A \otimes B \rightarrow T(A \otimes B)$ to be the composite

$$
T A \otimes B \xrightarrow{s_{T A, B}} B \otimes T A \xrightarrow{t_{B, A}} T(B \otimes A) \xrightarrow{T\left(s_{B, A}\right)} T(A \otimes B) .
$$

$T$ is commutative if the diagram

commutes for all objects $A$ and $B$ in $\mathcal{C}$.
If a monad is commutative then we get that its Kleisli category is symmetric monoidal for free with the monoidal tensor $\otimes$ (on objects) and unit being the same as in the underlying category $\mathcal{C}$.

Lemma 4.1.3 ([35]). If $T$ is a commutative monad over a symmetric monoidal category $\mathcal{C}$, then $\mathbf{K l}(T)$ is symmetric monoidal.

Commutative monads over Set also come with canonical copy/delete comonoid structures for every object. Copying $c_{X}: X \rightarrow T(X \times X)$ is given by

$$
X \xrightarrow{\Delta} X \times X \xrightarrow{\eta} T(X \times X)
$$

and deleting $d_{X}: X \rightarrow I$ is given by

$$
X \xrightarrow{!}\{\star\} \xrightarrow{\eta} T(\{\star\}) .
$$

From this comonoid structure we obtain canonical projections

$$
X \otimes Y \xrightarrow{\mathrm{id} \otimes d} X \otimes I \xrightarrow{\rho} X
$$

and

$$
X \otimes Y \xrightarrow{d \otimes \mathrm{id}} I \otimes Y \xrightarrow{\lambda} Y
$$

Crucially, it is not guaranteed that the monoidal tensor of $\mathbf{K l}(T)$ is cartesian.

### 4.2 The category of sets and random functions

We now turn to the category of interest for this section.
The finitary distribution monad $D$ : Set $\rightarrow$ Set maps a set $X$ to the set of finitary probability distributions on $X$ (finitary in the sense that only finitely many elements are assigned non-zero probability).

Definition 4.2.1 (Finitary distribution monad). Define $D:$ Set $\rightarrow$ Set by

$$
D(X)=\left\{\alpha: X \rightarrow[0,1] \mid \operatorname{supp}(\alpha)<\aleph_{0}, \sum_{x \in \operatorname{supp}(\alpha)} \alpha(x)=1\right\}
$$

where $\operatorname{supp}(\alpha)$ is $\{x \in X \mid \alpha(x) \neq 0\}$, the support of $\alpha$. $D$ acts on morphisms by

$$
D(f: X \rightarrow Y)(\alpha: D(X))(y)=\sum_{f(x)=y} \alpha(x)
$$

The monad structure of $D$ is given as follows. The unit is given by

$$
\eta_{X}(x)=\delta_{x}
$$

where

$$
\delta_{x}\left(x^{\prime}\right)= \begin{cases}1 & \text { if } x=x^{\prime} \\ 0 & \text { otherwise }\end{cases}
$$

and the extension $f^{\dagger}: D(X) \rightarrow D(Y)$ of $f: X \rightarrow D(Y)$ is

$$
f^{\dagger}(\alpha)=\sum_{x \in \operatorname{supp}(\alpha)} f(x)(y)
$$

Lemma 4.2.2. $D$ is a commutative monad.
Corollary 4.2.3. $\mathbf{K l}(D)$ is symmetric monoidal with canonical copy/delete comonoids and projection maps.

The monoidal tensor in $\mathbf{K l}(D)$, which represents joint distributions, is not cartesian.
Definition 4.2.4 (Copy/delete comonoid for $\mathbf{K l}(D)$ ). The copying map in $\mathbf{K l}(D)$ is given explicitly by $c_{X}: X \rightarrow D(X \times X)$ where

$$
c(x)\left(x_{1}, x_{2}\right)= \begin{cases}1 & \text { if } x=x_{1}=x_{2} \\ 0 & \text { otherwise }\end{cases}
$$

The monoidal unit of $\mathbf{K l}(D)$ is terminal, and hence the deleting map $d_{X}: X \rightarrow D(\{\star\})$ must be given by $d(x)(\star)=1$. As this map is unique, we refer to it as !.

Whenever we are working in $\mathbf{K l}(D)$ we denote $c_{X}$ and $!_{X}$ by

respectively.
The canonical projections in $\mathbf{K l}(D)$ correspond to taking marginals.
Definition 4.2.5 (Marginals). Let $p: D(X \times Y)$ be a joint distribution. The left marginal $p_{X}: D(X)$ is given by $p_{X}(x)=\sum_{y \in \operatorname{supp}(p(x,-))} p(x, y)$. The right marginal $p_{Y}: D(Y)$ is given similarly by $p_{Y}(y)=\sum_{x \in \operatorname{supp}(p(-, y))} p(x, y)$. As diagrams, these are given by

and

respectively.
An important operation on probability distributions is Bayesian updating where an agent has some prior distribution (initial belief), makes an observation, and then updates their prior to a new, posterior distribution.
Definition 4.2.6 (Update operator). Let $X$ and $\Theta$ be sets. We think of $\Theta$ as a type which an agent has a probabilistic belief about, and $X$ as a type that will be observed by an agent. Define the update operator $\mathcal{U}_{\Theta}: D(\Theta \times X) \rightarrow(X \rightarrow D(\Theta))$ by

$$
\mathcal{U}_{\Theta}(p)(x)(\vartheta)=\frac{p(\vartheta, x)}{\sum_{\left(\vartheta^{\prime}, x\right) \in \operatorname{supp}(p)} p\left(\vartheta^{\prime}, x\right)}
$$

We also write $p(\vartheta \mid x)$ for $\mathcal{U}_{\Theta}(p)(x)(\vartheta)$.

Lemma 4.2.7. The update operator is natural in $\Theta$. That is, the following diagram commutes for any $f: \Theta_{1} \rightarrow \Theta_{2}$ :


Proof. Let $p \in D\left(\Theta_{1} \times X\right), x \in X, \vartheta_{1} \in \Theta_{1}$, and $\vartheta_{2} \in \Theta_{2}$. The top of the square is given by

$$
\begin{align*}
\left(D(f) \circ \mathcal{U}_{\Theta_{1}}\right)(p)(x)\left(\vartheta_{2}\right) & =\sum_{f\left(\vartheta_{1}\right)=\vartheta_{2}} \mathcal{U}_{\Theta_{1}}(p)(x)\left(\vartheta_{1}\right) \\
& =\frac{\sum_{f\left(\vartheta_{1}\right)=\vartheta_{2}} p\left(\vartheta_{1}, x\right)}{\sum_{\vartheta_{1}^{\prime} \in \operatorname{supp}(p(-, x))} p\left(\vartheta_{1}^{\prime}, x\right)} .
\end{align*}
$$

The bottom of the square is given by

$$
\begin{align*}
\mathcal{U}_{\Theta_{2}}(D(f \times X)(p))(x)\left(\vartheta_{2}\right) & =\frac{D(f \times X)(p)\left(\vartheta_{2}, x\right)}{\sum_{\vartheta_{2}^{\prime}} D(f \times X)(p)\left(\vartheta_{2}^{\prime}, x\right)} \\
& =\frac{\sum_{f\left(\vartheta_{1}\right)=\vartheta_{2}} p\left(\vartheta_{1}, x\right)}{\sum_{\vartheta_{2}^{\prime}} \sum_{f\left(\vartheta_{1}^{\prime}\right)=\vartheta_{2}^{\prime}} p\left(\vartheta_{1}^{\prime}, x\right)} .
\end{align*}
$$

The result follows, noting that the denominators of $(\star)$ and ( $\star \star$ ) are equal.

### 4.3 Bayesian open games

Definition 4.3.1 (Bayesian open game). A Bayesian open game is a morphism in $\mathbf{G a m e}_{\mathbf{K l}(D)}$. Explicitly, a Bayesian open game $\mathcal{G}:(X, S) \rightarrow(Y, R)$ is given by

1. A set of strategies $\Sigma$,
2. A play function

$$
\mathrm{P}: \Sigma \rightarrow \int^{\Theta: \text { Set }}(X \rightarrow D(\Theta \times Y)) \times((\Theta \times R) \rightarrow D(S))
$$

3. A best-response function

$$
\text { B : } \mathbf{L e n s}_{\mathbf{K l}(D)}((I, R),(X, Y)) \rightarrow \operatorname{Rel}(\Sigma)
$$

We refer to atoms in the category of Bayesian open games as Bayesian atoms.
We unpack the definition of the play function to emphasise that, when we wish to actually specify a Bayesian open game, it is usually easier to specify $\mathrm{P}(\sigma)$ as the equivalence class of a pair of morphisms.

### 4.4 Bayesian agents

We will now define Bayesian agents which, as with concrete open games, have constant bestresponse functions. Bayesian agents capture the notion of rational agents that

1. Have a correct prior about the various types in a game;
2. Update this prior based on an observation to a new posterior;
3. Attempt to maximise their expected utility given their posterior.

Definition 4.4.1 (Bayesian agent). Let $X, Y$ be sets. The Bayesian agent $\mathcal{A}_{(X, Y)}:(X, I) \rightarrow$ $(Y, \mathbb{R})$ is the Bayesian atom given by

1. $\Sigma_{\mathcal{A}}=X \rightarrow D(Y)$,
2. 

$$
\mathrm{P} \sigma=\left[\sigma,!_{\mathbb{R}}\right]=\quad X \quad-\sigma \quad Y \quad \mathbb{R}
$$

3. The selection function $\varepsilon: \operatorname{Lens}_{\mathbf{K l}(D)}((I, \mathbb{R}),(X, Y)) \rightarrow \mathcal{P}(X \rightarrow D(Y))$ is given by

$$
\begin{aligned}
\varepsilon_{\mathcal{A}}([p, k])=\{\sigma: X \rightarrow D(Y) \mid & \forall x \in \operatorname{supp}\left(p_{2}\right) \\
& \left.\sigma(x) \in \underset{\alpha \in D(Y)}{\arg \max }\left(\mathbb{E}\left[k^{\dagger}\left(\mathcal{U}_{\Theta}(p)(x), \alpha\right)\right]\right)\right\} .
\end{aligned}
$$

It is worth explaining this last formula in words. $\mathcal{A}_{(X, Y)}$ represents an agent choosing an element of $Y$ after observing an element of $X$. The context of the decision consists of a set $\Theta$ of unobservable states, a prior joint distribution $p: D(\Theta \times X)$ on unobservable and observable states, and a utility function $k: \Theta \times Y \rightarrow D(\mathbb{R})$ that depends on the unobservable state and the agent's choice. The optimality condition says that for all observations $x$ that the agent could make with nonzero probability, the strategy $\sigma(x)$ maximises the expected value of $k\left(\vartheta^{\prime},-\right)$, where $\vartheta^{\prime}=\mathcal{U}_{\Theta}(p)(x)$ is the posterior distribution on the unobservable state given the observation of $x$.

Lemma 4.4.2. The selection function of a Bayesian agent is well-defined. That is, it is independent of the choice of representative of the coend equivalence relation.

Proof. This result follows from the fact that Bayesian updating is natural in the bound type of a coend lens (4.2.7).

In the next definition we formalise the idea that a player in a game might be assigned a (gametheoretic) type on which their utility function depends. We can do this simply using a Bayesian agent and a copying computation.

Definition 4.4.3. Let $\mathcal{A}_{(X, Y)}:(X, I) \rightarrow(Y, \mathbb{R})$ be a Bayesian agent. Define $\mathcal{A}_{(X, Y)}^{\Delta}:(X, I) \rightarrow$ $(X \times Y, \mathbb{R})$ to be the Bayesian open game $\left(\operatorname{id}_{(X, 1)} \otimes \mathcal{A}_{(X, Y)}\right) \circ \Delta_{X}$, or as a string diagram,


Lemma 4.4.4. $\mathcal{A}_{(X, Y)}^{\Delta}$ is explicitly given, up to isomorphism, by

1. $\Sigma_{\mathcal{A}_{(X, Y)}^{\Delta}}=X \rightarrow D(Y)$;
2. $\mathrm{P}_{\mathcal{A}_{(X, Y)}^{\Delta}}: \Sigma_{\mathcal{A}_{(X, Y)}^{\Delta}} \rightarrow \operatorname{Lens}_{\mathbf{K l}(D)}((X, I),(X \times Y, \mathbb{R}))$ is given by

3. Let $[p, k] \in \operatorname{Lens}_{\mathbf{K 1}(D)}((I, \mathbb{R}),(X, X \times Y))$. Best response is given, up to isomorphism, by $\mathrm{B}_{\mathcal{A}_{(X, Y)}^{\Delta}}([p, k])=\mathrm{B}_{\mathcal{A}_{(X, Y)}}\left([p, k]^{\prime}\right)$ where $[p, k]^{\prime}$ is the context given by


Proof. This result follows from definition chasing, noting that $\mathcal{A}_{(X, Y)}$ is the only component with non-trivial strategy profile set.

Lemma 4.4.5. Let $\mathcal{A}_{\left(X_{i}, Y_{i}\right)}$ be Bayesian agents for $i \in\{1, \cdots, n\}$. Then $\mathcal{A}^{\Delta^{n}}=\bigotimes_{i=1}^{n} \mathcal{A}_{\left(X_{i}, Y_{i}\right)}^{\Delta}$ is explicitly given as follows.

1. $\Sigma_{\mathcal{A}^{n}}=\prod_{i=1}^{n}\left(X_{i} \rightarrow D\left(Y_{i}\right)\right) ;$
2. The play function

$$
\mathrm{P}_{\mathcal{A}^{\Delta^{n}}}: \Sigma_{\mathcal{A}^{\Delta^{n}}} \rightarrow \operatorname{Lens}_{\mathbf{K l}(D)}\left(\left(\prod_{i=1}^{n} X_{i}, I\right),\left(\prod_{i=1}^{n}\left(X_{i} \times Y_{i}\right), \mathbb{R}^{n}\right)\right)
$$

is given by

3. Let $[p, k] \in \operatorname{Lens}_{\mathbf{K l}(D)}\left(\left(I, \mathbb{R}^{n}\right),\left(\prod_{i} X_{i}, \prod_{i} X_{i} \times Y_{i}\right)\right)$. Best response

$$
\mathrm{B}_{\mathcal{A}^{\Delta^{n}}}([p, k])\left(\sigma_{1}, \cdots, \sigma_{n}\right)
$$

is given, up to isomorphism, by

$$
\begin{aligned}
\prod_{i=1}^{n}\left\{\sigma_{i}: X_{i} \rightarrow D\left(Y_{i}\right) \mid \forall x_{i}\right. & \in \operatorname{supp}\left(p^{\sigma_{-i}}\right) \\
\sigma_{i}\left(x_{i}\right) & \in \underset{\alpha_{i} \in D\left(Y_{i}\right)}{\arg \max }\left(\mathbb{E}\left[k^{\sigma_{-i}}\left(\mathcal{U}_{\Theta_{-i}}\left(p^{\sigma_{-i}}\right)\left(x_{i}\right), \alpha_{i}\right]\right)\right\}
\end{aligned}
$$

where $\left[p^{\sigma_{-i}}, k^{\sigma_{-i}}\right]$ is the context given by


Proof. (1) and (2) follow easily from definitions. As for (3), we need to prove that the local context for each $\mathcal{A}_{\left(X_{i}, Y_{i}\right)}$ is $\left[p^{\sigma_{-i}}, k^{\sigma_{-i}}\right]$. Note that the previous Lemma 4.4.4 serves as the base case $(n=1)$ for an induction argument. The result then follows easily by considering that

$$
\mathrm{B}_{\mathcal{A}^{\Delta^{n}}}([p, k])\left(\sigma_{1}, \cdots, \sigma_{n}\right)=\mathrm{B}_{\mathcal{A}_{\left(x_{1}, Y_{1}\right)}^{\Delta}}\left([p, k]_{1}\right)\left(\sigma_{1}\right) \times \mathrm{B}_{\bigotimes_{i=2}^{n} \mathcal{A}_{\left(X_{i}, Y_{i}\right)}^{\Delta}}\left([p, k]_{-1}\right)\left(\sigma_{-1}\right)
$$

and applying the inductive hypothesis, where $[p, k]_{-1}$ is the context


### 4.5 Decisions under risk

In this section we introduce another type of situation involving a Bayesian agent that can be modelled using Bayesian open games.

A decision problem under risk is a decision problem for which one can sensibly assign probabilities to possible outcomes. A good example is roulette. When making a bet in roulette, you can calculate the likelihood of success and also your expected return on any bet. Decision problems under risk are generally represented by Bayesian open games constructed from computations and precisely one Bayesian agent. A simple subclass of decision problems under risk are represented by Bayesian open games of the form

in which an agent $\mathcal{A}$ attempts to maximise their outcome which is, in part, dependent on the type $Z$ which $\mathcal{A}$ does not observe.

We now give a fully worked out example of a Bayesian open game in which an agent has a prior, makes an observation, updates their prior as a consequence of that observation, and then makes a prediction based on their posterior.

Example 4.5.1 (Biased coin). Suppose we give an agent $\mathcal{A}$ a biased coin which lands on one side $75 \%$ of the time and the other side $25 \%$ of the time. It is not known which side the coin is biased towards, but it is known that it is equally likely to be biased towards heads as towards tails. $\mathcal{A}$ flips the coin whilst another identical coin (i.e. another coin biased the same way) is flipped in secret. $\mathcal{A}$ observes her coin flip and is then asked to predict which side up the secret coin landed. If she is correct she receives an outcome of 1 with probability 1 . If she is wrong she receives an outcome of 0 with probability 1 . The optimal strategy for $\mathcal{A}$ is to guess that the coin flipped in secret will land the same way up as the coin she flipped. If, for instance, $\mathcal{A}$ 's coin comes up heads, then there is a $75 \%$ chance that both coins are biased towards heads. Consequently the coin flipped in secret is more likely to show heads. A symmetric argument applies if $\mathcal{A}$ 's coin shows tails.

We can represent this game using the open game

where

$$
\begin{aligned}
p: D\left(\{H, T\}^{2}\right)= & \frac{1}{2}\left(\frac{9}{16}(T, T)+\frac{1}{16}(H, H)+\frac{3}{16}(T, H)+\frac{3}{16}(H, T)\right) \\
& +\frac{1}{2}\left(\frac{9}{16}(H, H)+\frac{1}{16}(T, T)+\frac{3}{16}(H, T)+\frac{3}{16}(T, H)\right) \\
= & \frac{5}{16}(H, H)+\frac{5}{16}(T, T)+\frac{3}{16}(T, H)+\frac{3}{16}(H, T)
\end{aligned}
$$

and

$$
\begin{gathered}
u:\{H, T\}^{2} \rightarrow D \mathbb{R} \\
(x, y)
\end{gathered}> \begin{cases}\delta_{1} & \text { if } x=y \\
\delta_{0} & \text { otherwise. }\end{cases}
$$

Explicitly, the game is given by $\mathcal{G}:=u \circ\left(\mathcal{A} \otimes \operatorname{id}_{\{H, T\}}\right) \circ p$. Note that $\Sigma_{\mathcal{G}} \cong \Sigma_{\mathcal{A}}=\{H, T\} \rightarrow$ $D(\{H, T\})$. Also note that there is precisely one context for $\mathcal{G}$ since its type is $I \rightarrow I$ (that is, $(1,1) \rightarrow(1,1))$ and, moreover, as the best-response functions for $u, i d_{\{H, T\}}$, and $p$ are trivial, the best-response function for $\mathcal{G}$ is isomorphic to the constant relation

$$
\left.\begin{array}{rl}
\mathrm{B}_{\mathcal{A}}([p, k])=\{\sigma:\{H, T\} \rightarrow D(\{H, T\}) \mid & \forall
\end{array}\right) \in\{H, T\}, 0 \text {. } \quad \begin{aligned}
& \\
& \left.\sigma(x) \in \arg \max _{\alpha \in D(\{H, T\})} \mathbb{E}\left[k\left(\mathcal{U}_{\{H, T\}}(p)(x), \alpha\right)\right]\right\} .
\end{aligned}
$$

The posterior $\mathcal{U}_{\{H, T\}}(p)(H)(H)$ is given by

$$
\begin{aligned}
\mathcal{U}_{\{H, T\}}(p)(H)(H) & =\frac{p(H, H)}{\sum_{(\vartheta, H) \in \operatorname{supp}(p)} p(\vartheta, H)} \\
& =\frac{5}{8}
\end{aligned}
$$

and hence $\mathcal{U}_{\{H, T\}}(p)(H)(T)=\frac{3}{8}$. Similarly, $\mathcal{U}_{\{H, T\}}(p)(T)=\frac{5}{8} T+\frac{3}{8} H$. It follows that

$$
\begin{aligned}
\arg \max _{\alpha}\left(\mathbb{E}\left[k^{\dagger}\left(\mathcal{U}_{\{H, T\}}(p)(H), \alpha\right)\right]\right) & =\left\{\delta_{H}\right\} \\
\arg \max _{\alpha}\left(\mathbb{E}\left[k^{\dagger}\left(\mathcal{U}_{\{H, T\}}(p)(T), \alpha\right)\right]\right) & =\left\{\delta_{T}\right\} .
\end{aligned}
$$

Hence $\mathrm{B}_{\mathcal{A}}([p, k])$ is a singleton set containing the strategy $\sigma$ where $\sigma(H)=\delta_{H}$ and $\sigma(T)=\delta_{T}$, as expected.

### 4.6 Relating Bayesian Open Games to Bayesian Games

The example given in the previous section illustrates two important improvements over the version of open games as introduced in [13]. There, everything is deterministic - the environment, the players moves etc. Here, the environment as well as the players' moves can be probabilistic.

Still, we only consider a single player. We could consider examples with more players and with probabilistic behaviour induced by nature or games where probabilistic behaviour is key, such as matching pennies. In [15] a definition of open games is given that can handle mixed strategies.

Instead, we want to focus on another aspect of Bayesian open games, which we believe is much less obvious, particularly to readers with no game theory background: From a game-theoretic perspective, we introduced a new solution concept. For random environments and probabilistic moves, there is no real difference to standard Nash equilibria. However, there are important classes of games for which the new equilibrium notion brings material change.

Reconsider Example 4.5.1. Here, the agent has some prior information about the nature of the coin. Given observations from the coin and the prior information, the agent updates beliefs about the true nature of the coin and maximises accordingly. Thus, we maintain not only an assumption that the agent maximises but also how she deals with information.

The example is very simple. After all, there is only one agent. Our interest lies of course in the interaction of such Bayesian agents. While the strategic reasoning is more complicated in such situations the overall structure is the same as in Example 4.5.1: Each agent has access to some prior information, some (partial) information is revealed, and the agent makes a decision.

In the game-theory literature such a game is called a 'Bayesian game' or also 'game in Bayesian form' and the solution concept we sketched is called 'Bayesian Nash Equilibrium'. Next, we introduce the standard economic notion of such games, then provide a simple example, an auction, and followed by a general construction that allows us to translate an arbitrary Bayesian game into a Bayesian open game.

### 4.6.1 Games in Bayesian form

Denote with $\Gamma^{B}=\left(N,\left(A_{i}\right)_{i=1}^{N},\left(\Theta_{i}\right)_{i=1}^{N}, \pi,\left(u_{i}\right)_{i=1}^{N}\right)$ a Bayesian game. Before we explain the components we add some useful notation: we write $\Theta_{-i}=\prod_{j \in N, j \neq i} \Theta_{j}$ for the possible combinations of $\Theta_{j}$ other than $i$.

- $N$ denotes the number of players.
- $A_{i}$ defines a set of actions available to the players. ${ }^{3}$
- $\Theta_{i}$ defines the a set of (player) types each player $i$ can have.

[^3]- $\pi: D\left(\prod_{i=1}^{N} \Theta_{i}\right)$ is a joint prior distribution on the types that is common knowledge.
- $u_{i}: \prod_{i=1}^{N} \Theta_{i} \times \prod_{i=1}^{N} A_{i} \rightarrow \mathbb{R}$ defines a utility for any profile of player types and any profile of actions.

This definition is rather cryptic at first. So let us try to dissect its components. Then we will compare it to the definition of a normal form which we introduced in 2.12.4. This hopefully further helps to make sense of it.

Let us begin with $\Theta_{i}$, the player type, which is probably most obscure and most central at the same time. One way to think of types is imagining player $i$ having different realisations or versions, one of which ends up playing the game. Each realisation, or type, summarises information relevant for the game. Often, this information will concern the payoffs. ${ }^{4}$ For example, player $i$ may consider to bid in an auction. His type determines his evaluation for the good, that is, how much he values owning the good. It is easily conceivable that $i$ may have different evaluations before the auction really takes place.

This example also illustrates another aspect of this structure. While player $i$ may have different types, before he has to choose his action, he typically will know his type. So, once the bidding starts, he knows how much he values the good. But crucially, he does not know the types of the other players, only a probability distribution - conditional on his own observed type. The last aspect also already alludes to the role of Bayesian updating: Given his own observed type, a player may refine his belief about the other players' types.

The payoff function $u_{i}$, as with a normal form, maps choices of players into a payoff. Here, however, player $i$ 's type may affect the payoff. Think again about an auction, naturally the bidding behaviour will affect $i$ 's payoff. But so does his type, the evaluation for the good. Note that the definition above includes the possibility that the whole realisation of types affects $i$ 's payoff. In the case of bidding for a private good, only $i$ 's type will typically be relevant.

Lastly, $A_{i}$ represents the set of actions player $i$ can take in the game. Note, that we intentionally refer to this as actions and not as strategies. To understand this it is best to compare the above definition to a normal-form game.

Recall its notation: $\Gamma=\left(N,\left(S_{i}\right)_{i=1}^{N},\left(u_{i}\right)_{i=1}^{N}\right)$. Here, $S_{i}$ refers to a strategy which is a complete contingent plan for all possible occasions where $i$ can make a move. Note that $S_{i}$ can be a shorthand for some complicated dynamic structure where $i$ moves several times.

Analogously, $A_{i}$ in the Bayesian game refers to a complete contingent plan once the game 'begins', i.e. after $i$ has learned his type. This could involve a complicated dynamic structure. Crucially, however, this is not the same as a strategy for the Bayesian game. Why? A strategy of the Bayesian game must include a contingent plan also for each realisation of $i$ 's type!

Formally, a pure strategy for a game in Bayesian form is a mapping $s_{i}: \Theta_{i} \rightarrow A_{i}$ for all $\Theta_{i}$. A behavioural strategy $\sigma_{i}$ for player $i$ is a mapping $\sigma_{i}: \Theta_{i} \rightarrow D\left(A_{i}\right)$, i.e for each type of player $i$ a behavioural strategy assigns a probability distribution over the available actions. ${ }^{5}$

For each player we can now define the (conditional) expected utility for player $i$ given behavioural strategy profile $\sigma$ with

$$
E u_{i}\left(\sigma \mid t_{i}\right) \equiv \sum_{t_{-i} \in \Theta_{-i}} p_{\pi}\left(t_{-i} \mid t_{i}\right) u_{i}\left(\left(t_{i}, t_{-i}\right), \sigma\right)
$$

Note, that $p_{\pi}(\cdot \mid \cdot)$ updates the prior information $\pi$ for player $i$ given Bayes' rule. Equipped with all that we can finally define a Bayesian Nash equilibrium.

[^4]Definition 4.6.2 (Bayesian equilibrium). A (behavioural) strategy profile $\sigma^{*}=\left(\sigma_{1}^{*}, \sigma_{2}^{*}, \ldots, \sigma_{N}^{*}\right)$ is a Bayesian Nash equilibrium if for each player $i \in N$, each type $\vartheta_{i} \in \Theta_{i}$, and each possible action $a_{i} \in A_{i}$ it holds that:

$$
E u_{i}\left(\sigma^{*} \mid \vartheta_{i}\right) \geq E u_{i}\left(a_{i}, \sigma_{-i} \mid \vartheta_{i}\right)
$$

### 4.6.3 An auction example

Suppose two agents are bidding for a good that is being sold in a first-price sealed-bid auction. Here 'first-price' means that whichever agent bids higher receives the good and pays her/his own bid, with the other agent neither gaining nor losing anything. 'Sealed-bid' means neither player can observe the other's bid, meaning the bidding is effectively simultaneous.

Both players have a private valuation for the good, which are drawn from a joint random distribution $\pi: D(\Theta \times \Theta)$, where $\Theta \subseteq[0, \infty)$ is the range of possible values. As usual, we assume this prior is common knowledge. One consequence thereof is that a player knowing their own valuation can update their beliefs about the other's valuation. Obviously, whether there is something to learn from one's own valuation about others' valuations depends on the type of good being auctioned. For instance, if we are bidding on a construction contract, then the value that contract has for me will be correlated with your evaluation. If, however, we are bidding on a painting by some obscure painter, it is less clear that I can learn something about others' valuations.

The winning bidder's utility is given by their valuation of the good minus the bid that they must pay. (This can be negative if the bid is higher than the valuation.) We must also choose a way to resolve a tie if the bids are equal: we assume that the winner is determined by a fair coin flip. Thus the expected utilities, for private valuations $\vartheta_{i} \in \Theta$ and bids $b_{i} \in \Theta$, are given by

$$
U\left(\vartheta_{1}, b_{1}, \vartheta_{2}, b_{2}\right)= \begin{cases}\left(\vartheta_{1}-b_{1}, 0\right) & \text { if } b_{1}>b_{2} \\ \left(0, \vartheta_{2}-b_{2}\right) & \text { if } b_{1}<b_{2} \\ \left(\frac{\vartheta_{1}-b_{1}}{2}, \frac{\vartheta_{2}-b_{2}}{2}\right) & \text { if } b_{1}=b_{2}\end{cases}
$$

In order to formalise this situation as a Bayesian game as defined in the previous section, we take $N=2, A_{1}=A_{2}=\Theta$, and $\pi$.

The Bayesian open game describing this situation is given by the string diagram


As a morphism $\mathcal{G}: I \rightarrow I$ in the category of Bayesian open games, this has $\Sigma_{\mathcal{G}}=(\Theta \rightarrow$ $D(\Theta)) \times(\Theta \rightarrow D(\Theta))$, the set of behavioural strategy profiles for the auction, and $B_{\mathcal{G}}(*)$ the best-response relation for behavioural strategy profiles.

### 4.6.4 General construction for games in Bayesian form

Generalising the previous example, we can give a construction for converting any game in Bayesian form into a Bayesian open game of type $I \rightarrow I$ that has the same best-response relation.

Consider a general Bayesian game $\left(N,\left(A_{i}\right),\left(\Theta_{i}\right), \pi, U\right)$. For each player $i$, recall the copying agent $\mathcal{A}_{\left(\Theta_{i}, A_{i}\right)}^{\Delta}:\left(\Theta_{i}, 1\right) \rightarrow\left(\Theta_{i} \times A_{i}, \mathbb{R}\right)$. Their tensor product has type

$$
\bigotimes_{i=1}^{N} \mathcal{A}_{\left(\Theta_{i}, A_{i}\right)}^{\Delta}:\left(\prod_{i=1}^{N} \Theta_{i}, 1\right) \rightarrow\left(\prod_{i=1}^{N}\left(\Theta_{i} \times A_{i}\right), \mathbb{R}^{N}\right)
$$

The full game is given by

$$
I \xrightarrow{\langle\pi, 1\rangle}\left(\prod_{i=1}^{N} \Theta_{i}, 1\right) \xrightarrow{\left.\bigotimes_{i=1}^{N} \mathcal{A}_{(\Theta}^{\Delta} \Theta_{i}, A_{i}\right)}\left(\prod_{i=1}^{N}\left(\Theta_{i} \times A_{i}\right), \mathbb{R}^{N}\right) \xrightarrow{\left\langle U, \mathrm{id}_{\mathbb{R}^{N}}\right\rangle}\left(\mathbb{R}^{N}, \mathbb{R}^{N}\right) \xrightarrow{\varepsilon_{\mathbb{R}^{N}}} I .
$$

Proposition 4.6.5. Let $G=\left(N,\left(A_{i}\right),\left(T_{i}\right), \pi, U\right)$ be a game Bayesian form, and let

$$
c=(1, \pi, U): \mathbb{C}\left(\left(\prod_{i=1}^{N} T_{i}, 1\right),\left(\prod_{i=1}^{N}\left(T_{i} \times A_{i}\right), \mathbb{R}\right)\right) .
$$

Then $B^{\bigotimes_{i=1}^{N} \mathcal{A}_{\left(T_{i}, A_{i}\right)}^{\Delta}}(c)$ is the best-response relation for behavioural strategy profiles of $G$.
The main part of the proof of this result is contained in Lemma 4.4.5, which characterises the tensor product of copying agents.

### 4.7 Extensive form

The normal-form representation as well as the representation of Bayesian Games really are strategycentric. They both condense the strategic situation in a way that makes them most amenable to analysis from the point of view of a specific equilibrium concept such as Nash/Bayesian Nash equilibrium.

Open games have baked in a notion of equilibrium and thus certainly share this focus. And for that reason, comparisons to classical game theory in this paper have been limited to these representations.

In classical game theory, an alternative to the normal-form representation of a game is its extensive form. This representation can also be used for games in Bayesian form. We will not give a definition here, as standard definitions are cumbersome and will not add much value. What is important though is what the extensive form represents: it provides a detailed account about which player knows what at which point and what actions he has then available. One can think of the extensive form as a detailed account of how a game and the players' information unfolds. Thus, it provides much more information than its corresponding normal form, to which each extensive form can be reduced.

For instance, one reason to consider the extensive form representation of a game is that it explicitly captures temporal information, say, if one player moves before the other player. ${ }^{6}$

Open games as a representation of a strategic situation have much in common with its extensive form representation. After all they provide an expressive language to represent a detailed account of a strategic interaction. In the following, we will make the relationship between Bayesian open games as presented in this paper and the extensive form representation of Bayesian games more concrete.

We begin with a concrete example that illustrates how open games deal with sequential moves of players in the presence of asymmetric information.

### 4.7.1 The market for lemons

The following is a simple version of a classic model going back to [1]. ${ }^{7}$ There are two players, a seller and a buyer. The seller wants to sell a used car to the buyer; he offers the buyer a price

[^5]$P$; the buyer can either accept or reject the offer. Both players have some evaluation for the car which will affect their profit/utility.

Their evaluation depends on the quality $Q$ of the car. It is drawn from some fixed distribution $\pi$. Critically, the seller observes the quality before offering a price to the buyer. The buyer, however, does not directly observe the quality.

The open game below represents this situation.


As Akerlof [1] observed, due to the asymmetry in information, a transaction may not occur even though such a transaction would be efficient (and would take place if both players were perfectly informed). The reason is simple. For each price the buyer has to take into consideration what quality of car the seller would be willing to offer. As each seller will only offer prices equal to or above his evaluation for the car, a low price indicates low quality. However, a high price does not necessarily reflect high quality; it could just mean that the seller is charging a high price for an actual 'lemon'. Thus, the buyer, given his prior will update his estimate of the quality for each price offered. And depending on the parameters, transactions may only occur for low quality cars but not for high quality cars.

Note that, similar to our remark in 2.12.3, even though the interaction (and the representation as an open game) is sequential, the equilibrium notion is static and does not consider the timing of moves. ${ }^{8}$

### 4.7.2 A general construction

To translate an arbitrary extensive form game to a string diagram requires an additional external choice operator like the one developed in [22], which is currently not well understood in the Bayesian setting. We will instead give a much simpler translation for a certain subclass of extensive form games, roughly those whose trees are perfectly balanced. Specifically we require:

- All paths from the root to a leaf have the same length
- All information sets intersect only one level of the tree (where the $i$ th level is the set of all nodes that are $i$ steps away from the root)
- All information sets on the same level are owned by the same player (the root node is always owned by Nature)
- All nodes on the same level have the same number of outgoing edges, labelled by the same edge labels

Any extensive form game can be converted into one satisfying these conditions that is equivalent for most game-theoretic purposes, by adding (potentially many) dummy nodes and moves leading to redundant subtrees with only strictly dominated payoffs. Obviously this is unsatisfactory, and the construction in this section should be thought of as an argument that Bayesian open games have the same 'expressive power' as extensive form, and not as a practical translation.

[^6]Consider an extensive form game $G$ satisfying these conditions, with $N>0$ players, and $L>0$ 'ordinary' levels with choice nodes, together with level 0 (the root) and level $L+1$ (the leaves).

At an ordinary level $1 \leq i \leq L$, suppose the set of nodes is $X_{i}$, the information equivalence relation is $\sim_{i}$ (and so the set of information sets at level $i$ is $X_{i} / \sim_{i}$ ), all information sets are owned by player $1 \leq P_{i} \leq N$, and the set of outgoing edge labels for all nodes is $A_{i}$. We summarise the connectivity between level $i$ and level $i+1$ as a transition function $q_{i}: X_{i} \times A_{i} \rightarrow X_{i+1}$.

The data at the root node is a probability distribution $\pi: D\left(X_{1}\right)$ on the Nature moves. Finally, $X_{L+1}$ is the set of leaf nodes, and the assignment of payoffs at leaf nodes is summarised as a utility function $U: X_{L+1} \rightarrow \mathbb{R}^{N}$.

We translate each level of the tree into a Bayesian open game, which we then sequentially compose together. An ordinary level $1 \leq i \leq L$ will be translated to an open game of type $\mathcal{G}_{i}:\left(X_{i}, \mathbb{R}^{N}\right) \rightarrow\left(X_{i+1}, \mathbb{R}^{N}\right)$, level 0 to an open game $\mathcal{G}_{0}:(1,1) \rightarrow\left(X_{1}, \mathbb{R}^{N}\right)$ and level $L+1$ to an open game $\mathcal{G}_{L+1}:\left(X_{L+1}, \mathbb{R}^{N}\right) \rightarrow(1,1)$. Thus, the final result will be an open game $\mathcal{G}=\mathcal{G}_{L+1} \circ \cdots \circ \mathcal{G}_{0}:(1,1) \rightarrow(1,1)$.

We define $\mathcal{G}_{0}$ to be the computation $\langle\pi,!\rangle$, where $\pi: 1 \rightarrow D\left(X_{1}\right)$ is the prior considered as a state in the Kleisli category, and $!: \mathbb{R}^{N} \rightarrow D(1)=1$ is the discard map. That is to say, it is the atomic open game with $\Sigma_{\mathcal{G}_{0}}=\{[\pi,!]\}$ and $B_{\mathcal{G}_{0}}(c)=\{[\pi,!]\}$ for all contexts $c \in \mathbb{C}\left((I, I),\left(X_{1}, \mathbb{R}^{N}\right)\right)$. $\mathcal{G}_{0}$ is simply depicted by the string diagram

$O \longleftarrow \mathbb{R}^{N}$
and the optic $[\pi,!]$ by the coend diagram


Similarly, we define $\mathcal{G}_{L+1}$ to be the atomic open game with $\Sigma_{\mathcal{G}_{L+1}}=\{U\}$ and $B_{\mathcal{G}_{L+1}}(c)=\{U\}$ for all contexts $c \in \mathbb{C}\left(\left(X_{L+1}, \mathbb{R}^{N}\right),(I, I)\right)$, where $U:\left(X_{L+1}, \mathbb{R}^{N}\right) \rightarrow(I, I)$ is the optic given by $\left(\mathbb{R}^{N}, U, \mathrm{id}\right)$, with the coend diagram

$\mathcal{G}_{L+1}$ is itself denoted by the string diagram


The open game $\mathcal{G}_{i}$ for $1 \leq i \leq L$ is defined by the string diagram


Unwinding this definition, we get that

- $\Sigma_{\mathcal{G}_{i}} \cong \Sigma_{\mathcal{A}_{\left(X_{i} / \sim_{i}, A_{i}\right)}}=X_{i} / \sim_{i} \rightarrow D\left(A_{i}\right)$
- For $\sigma_{i}: X_{i} / \sim_{i} \rightarrow D\left(A_{i}\right), P_{\mathcal{G}_{i}}\left(\sigma_{i}\right)$ is the optic $\left(X_{i}, \mathbb{R}^{n}\right) \rightarrow\left(X_{i+1}, \mathbb{R}^{n}\right)$ given by the coend diagram

- For a context $c=(\Theta, p, k):\left(I, \mathbb{R}^{N}\right) \rightarrow\left(X_{i}, X_{i+1}\right), B_{\mathcal{G}_{i}}(c)$ is the constant relation

$$
\begin{aligned}
\left\{\sigma_{i}: X_{i} / \sim_{i} \rightarrow D\left(A_{i}\right) \mid\right. & \sigma_{i}\left(\left[x_{i}\right]_{\sim_{i}}\right) \in \arg \max _{a_{i} \in A_{i}} \mathbb{E}\left[\pi_{P_{i}}\left(k\left(p_{1}, q\left(p_{2}, a_{i}\right)\right)\right) \mid p_{2}=x_{i}\right] \\
& \text { for all } \left.x_{i} \in \operatorname{supp}\left(p_{2}\right)\right\},
\end{aligned}
$$

where $\mathbb{E}\left[\pi_{P_{i}}\left(k\left(p_{1}, q\left(p_{2}, a_{i}\right)\right)\right) \mid p_{2}=x_{i}\right]$ is the conditional expectation of

$$
I \xrightarrow{p} \Theta \otimes X_{i} \xrightarrow{\Theta \otimes q\left(-, a_{i}\right)} \Theta \otimes X_{i+1} \xrightarrow{k} \mathbb{R}^{N} \xrightarrow{\pi_{P_{i}}} \mathbb{R}
$$

given that the right marginal of $p$ is $x_{i}$.
In order to characterise $\mathcal{G}=\mathcal{G}_{L+1} \circ \cdots \circ \mathcal{G}_{0}$ inductively we have a choice to either work forwards from $\mathcal{G}_{0}$ or backwards from $\mathcal{G}_{L+1}$, corresponding roughly to forward induction and backward induction in game theory. We choose the former. The inductive hypothesis is encapsulated in the following lemma.

Lemma 4.7.3. For $1 \leq i \leq L$, consider the extensive form game $G_{i}$ without specified payoffs given by truncating the original extensive form game $G$ to the first $i+1$ levels ( $i$ proper levels plus the root node). Note that the set of leaves of $G_{i}$ is (labelled by) $X_{i+1}$. Then the open game $\mathcal{G}_{i} \circ \cdots \circ \mathcal{G}_{0}:(1,1) \rightarrow\left(X_{i+1}, \mathbb{R}^{N}\right)$ is given by the following:

- $\Sigma_{\mathcal{G}_{i} \circ \ldots \circ \mathcal{G}_{0}}$ is the set of behavioural strategy profiles of $G_{i}$, namely $\Sigma_{\mathcal{G}_{i} \circ \ldots \circ \mathcal{G}_{0}}=\prod_{j=1}^{i}\left(X_{j} / \sim_{j} \rightarrow\right.$ $\left.D\left(A_{j}\right)\right)$
- $\left(\mathcal{G}_{i} \circ \cdots \circ \mathcal{G}_{0}\right)(\sigma)$ is the optic $(1, a,!)$, where $a: 1 \rightarrow D\left(X_{i+1}\right)$ is the probability distribution on the leaf nodes of $G_{i}$ resulting from playing the behavioural strategy profile $\sigma$
- Every context of $\mathcal{G}_{i} \circ \cdots \circ \mathcal{G}_{0}$ has the form

and thus is equivalent to one of the form $(1, *, k)$ where $*$ is the unique distribution on one point and $k: X_{i} \rightarrow D\left(\mathbb{R}^{N}\right)$, by taking $k$ to be the partial application of $k^{\prime}: \Theta \times X_{i+1} \rightarrow D\left(\mathbb{R}^{N}\right)$ to $p: D(\Theta)$. Then $\mathrm{B}_{\mathcal{G}_{i} \circ \cdots \circ \mathcal{G}_{0}}([*, k])(\sigma)$ is the set of best responses to $\sigma$ in the extensive form game given by $G_{i}$ together with payoffs, where the ith player's payoff from the leaf node $x \in X_{i}$ is $\mathbb{E}\left[k(x)_{i}\right]$ (where $k(x)_{i}$ is the $i$ th marginal of the joint distribution $k(x): D\left(\mathbb{R}^{N}\right)$ ).

By taking $i=L$ in the previous lemma and then post-composing with $\mathcal{G}_{L+1}$, we obtain the final result.

Theorem 4.7.4. $\mathcal{G}=\mathcal{G}_{L+1} \circ \cdots \circ \mathcal{G}_{0}:(1,1) \rightarrow(1,1)$ is given by the following:

- The set of strategy profiles of $\mathcal{G}$ is the set of behavioural strategy profiles of the original extensive form game $G$
- Every context is equivalent to one of the form $(1, *, *)$, and for a behavioural strategy profile $\sigma, \mathrm{B}_{\mathcal{G}}(1, *, *)(\sigma)$ is the set of best responses to $\sigma$ in $G$


## 5 Conclusion

In this last section we will discuss the wider context of this work, as well as some work in progress and future work.

Computer support. Due to the complexity of the definition of Bayesian open games, in practice computer support is necessary to work with real models in the framework. We have created such a software tool in the form of a Haskell library (available at https://github.com/jules-hedges/open-g ames-hs), consisting of a 'core' implementation of the monoidal category of Bayesian open games, together with a domain-specific embedded programming language that provides a higher level of description roughly equivalent to string diagrams. We have found this to be a practical tool for modelling, especially for rapid prototyping of models, in a variety of applied domains including auction design, governance modelling and blockchain protocol modelling. The implementation of this tool, and these applied case studies, will be described in a series of future papers.

Non-finitary probability. In Section 3 we gave a general theory of open games over a monoidal category. And then, in Section 4 specialised it to the Kleisli category of the finite support probability monad. This restriction to finite support distributions was done for simplicity, but the general machinery we have introduced is applicable to any category of probabilistic functions. Examples include the Kleisli categories of the Giry monads on measurable and Polish spaces [16], the Radon monad on compact Hausdorff spaces [38], and the Kantorovich monad on complete metric spaces [11]. In fact, we believe that Bayesian open games can be formulated over any Markov category [10], which encompasses all of these examples.

Other solution concepts. In this paper we have focused on the solution concept of Bayesian Nash equilibrium. It is one of the central solution concepts in applied economic modelling as it allows us to model situations of asymmetric information. It is probably no exaggeration that for most economic situations information is asymmetric and thus the default. Hence, this paper is an important step in making the theory of open games practically useful for a wide range of situations. As a side effect, both ordinary mixed strategy Nash equilibria and correlated equilibria are obtained as special cases of Bayesian Nash equilibrium. However, for dynamic games it is common to use equilibrium refinements such as perfect Bayesian equilibrium or sequential equilibrium, in which individual players' beliefs are represented explicitly. ${ }^{9}$ It is an open question whether it is possible to extend compositional game theory with these stronger solution concepts.

Behavioural aspects. A currently-unexplored benefit of Bayesian open games is that neither maximisation of real numbers nor Bayes' law is involved in the definition of the categorical structure, and instead both enter when decisions are defined. This means that the same framework can equally well accommodate agents who neither perfectly maximise, nor update their beliefs strictly according to Bayes' law. This could make our framework useful for behavioural game theory, where such alternatives are considered (see, e.g. [4]).

[^7]
### 5.1 Acknowledgements

The text of this paper is closely based on the first author's PhD thesis, [3]. The authors gratefully thank Mitchell Riley, whose suggestion to apply coend lenses to open games was crucial and very timely.

## References

[1] George A. Akerlof. The market for 'lemons': Quality uncertainty and the market mechanism. Quarterly Journal of Economics, 84(3):488-500, 1970. https://doi.org/10.2307/1879431. 40, 41
[2] Jeremy Avigad and Solomon Feferman. Gödel's functional ("Dialectica") interpretation. In Handbook of proof theory, volume 137 of Studies in logic and the foundations of mathematics, pages 337-405. Elsevier, 1998. https://doi.org/10.2307/420972. 3
[3] Joe Bolt. Probability and nondeterminism in compositional game theory. PhD thesis, University of Oxford, 2019. 11, 45
[4] Colin F Camerer. Behavioral game theory: Experiments in strategic interaction. Princeton University Press, 2011. 44
[5] Valeria de Paiva. The dialectica categories. Technical report, University of Cambridge, 1991. 3
[6] Martin Escardó and Paulo Oliva. Selection functions, bar recursion and backward induction. Mathematical Structures in Computer Science, 20(2):127-168, 2010. https://doi.org/10.1017/S0960129509990351. 7
[7] Brendan Fong. The algebra of open and interconnected systems. PhD thesis, University of Oxford, 2016. 1
[8] Brendan Fong, David Spivak, and Rémy Tuyéras. Backprop as functor: A compositional perspective on supervised learning. In Proceedings of Logic in Computer Science (LiCS) 2019. ACM, 2019. https://doi.org/10.1109/LICS.2019.8785665. 11
[9] Nate Foster, Michael Greenwald, Jonathan Moore, Benjamin Pierce, and Alan Schmitt. Combinators for bi-directional tree transformations: A linguistic approach to the view update problem. ACM Transactions on Programming Languages and Systems, 29(3), 2007. https://doi.org/10.1145/1232420.1232424. 3
[10] Tobias Fritz. A synthetic approach to Markov kernels, conditional independence and theorems on sufficient statistics. Advances in Mathematics, 370, 2020. https://doi.org/10.1016/j.aim.2020.107239. 44
[11] Tobias Fritz and Paolo Perrone. A probability monad as the colimit of spaces of finite samples. Theory and Applications of Categories, 34(7), 2019. 44
[12] Drew Fudenberg and Jean Tirole. Game theory. MIT Press, 1991. 1
[13] Neil Ghani, Jules Hedges, Viktor Winschel, and Philipp Zahn. Compositional game theory. In Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, pages 472-481. ACM, 2018. https://doi.org/10.1145/3209108.3209165. 1, 2, 37
[14] Neil Ghani, Clemens Kupke, Alasdair Lambert, and Fredrik Nordvall Forsberg. A compositional treatment of iterated open games. Theoretical Computer Science, 741:48-57, 2018. ISSN 0304-3975. https://doi.org/10.1016/j.tcs.2018.05.026. An Observant Mind : Essays Dedicated to Don Sannella on the Occasion of his 60th Birthday. 14
[15] Neil Ghani, Clemens Kupke, Alasdair Lambert, and Fredrik Nordvall Forsberg. Compositional game theory with mixed strategies: Probabilistic open games using a distributive law. In Proceedings of ACT 2019, volume 323 of Electronic Proceedings in Theoretical Computer Science, 2020. https://doi.org/10.4204/EPTCS.323.7. 2, 37
[16] Michèlle Giry. A categorical approach to probability theory. Categorical aspects of topology and analysis, pages 68-85, 1982. https://doi.org/10.1007/BFb0092872. 44
[17] John C Harsanyi. Games with incomplete information played by "Bayesian" players, I-III Part I. The basic model. Management Science, 14(3):159-182, 1967. https://doi.org/10.1287/mnsc.14.3.159. 2
[18] John C Harsanyi. Games with incomplete information played by "Bayesian" players part II. Bayesian equilibrium points. Management Science, 14(5):320-334, 1968. https://doi.org/10.1287/mnsc.14.5.320.
[19] John C Harsanyi. Games with incomplete information played by "Bayesian" players, part III. The basic probability distribution of the game. Management Science, 14(7):486-502, 1968. https://doi.org/10.1287/mnsc.14.7.486. 2
[20] Jules Hedges. Towards compositional game theory. PhD thesis, Queen Mary University of London, 2016. 13
[21] Jules Hedges. Coherence for lenses and open games. arXiv:1704.02230, 2017. https://doi.org/10.48550/arXiv.1704.02230. 7
[22] Jules Hedges. Morphisms of open games. Electronic Notes in Theoretical Computer Science, 341:151 - 177, 2018. ISSN 1571-0661. https://doi.org/10.1016/j.entcs.2018.11.008. Proceedings of the Thirty-Fourth Conference on the Mathematical Foundations of Programming Semantics (MFPS XXXIV). 11, 41
[23] Jules Hedges, Paulo Oliva, Evguenia Shprits, Viktor Winschel, and Philipp Zahn. Higher-order decision theory. In International Conference on Algorithmic DecisionTheory, pages 241-254. Springer, 2017. https://doi.org/10.1007/978-3-319-67504-6'17. 8
[24] Jules Hedges, Paulo Oliva, Evguenia Shprits, Viktor Winschel, and Philipp Zahn. Selection equilibria of higher-order games. In International Symposium on Practical Aspects of Declarative Languages, pages 136-151. Springer, 2017. https://doi.org/10.1007/978-3-319-51676-9"9. 7, 8
[25] Julian Hedges. Lenses for philosophers, 2018. Available at https://julesh.com/2018/08/16/le nses-for-philosophers/. 3
[26] Martin Hofmann and Benjamin Pierce. Positive subtyping. Information and Computation, 126:11-33, 1996. https://doi.org/10.1006/inco.1996.0031. 3
[27] F. Loregian. Coend Calculus, volume 468 of London Mathematical Society Lecture Note Series. Cambridge University Press, first edition, 07 2021. https://doi.org/10.1017/9781108778657. 2, 15, 27
[28] Andreu Mas-Colell, Michael D. Whinston, and Jerry Green. Microeconomic Theory. Oxford University Press, 1995. ISBN 0195102681. 40, 41, 44
[29] Michael Maschler, Eilon Solan, and Shmuel Zamir. Game Theory. Cambridge University Press, 2013. https://doi.org/10.1017/CBO9780511794216. 38, 44
[30] John Nash. Equilibrium points in n-person games. Proceedings of the national academy of sciences, 36(1):48-49, 1950. https://doi.org/10.1073/pnas.36.1.48. 1
[31] John Von Neumann and Oskar Morgenstern. Theory of games and economic behavior. Princeton university press, 1944. 1, 2
[32] Frank Oles. A category-theoretic approach to the semantics of programming languages. PhD thesis, Syracuse University, 1982. 3
[33] Martin J Osborne and Ariel Rubinstein. A course in game theory. MIT press, 1994. 1
[34] Luke Palmer. Making Haskell nicer for game programming. Available at https://lukepalmer .wordpress.com/2007/07/26/making-haskell-nicer-for-game-programming/, 2007. 3
[35] John Power and Edmund Robinson. Premonoidal categories and notions of computation. Mathematical Structures in Computer Science, 7(5):453-468, 1997. https://doi.org/10.1017/S0960129597002375. 30
[36] Mitchell Riley. Categories of optics. arXiv:1809.00738v2, 2018. 2, 3, 16, 17, 27, 29
[37] Mario Román. Open diagrams via coend calculus. In Proceedings of Applied Category Theory 2020, volume 333 of Electronic Proceedings in Theoretical Computer Science, 2021. https://doi.org/10.4204/EPTCS.333.5. 17
[38] T. Świrszcz. Monadic functors and convexity. Bulletin de l'Acaédemie Polonaise des Sciences, 22(1), 1974. 44
[39] Category Theory Zulip. Change of domain for coends. Archived at https://mattecapu.gith ub.io/ct-zulip-archive/stream/229199-learning:-questions/topic/topic_.22change.20of.20do main.22.20for.20coends.html, 2021. 27


[^0]:    Jules Hedges: jules.hedges@strath.ac.uk
    Philipp Zahn: philipp@20squares.xyz

[^1]:    ${ }^{1}$ Selection functions have been studied in a game-theoretic context in other form. See, e.g. [6] or [24].

[^2]:    ${ }^{2}$ In [36], lenses over a Kleisli category are called effectful optics.

[^3]:    ${ }^{3}$ Note, we intenionally refer to this as actions and not as strategies. This will be clear once we compare the Bayesian game to a normal-form game.

[^4]:    ${ }^{4}$ The concept of a type is much more expressive. It can express all sorts of uncertainty with respect to a game. We will not discuss this here as it would lead us too far astray. Compare, for instance, Chapters 9-11 in [29] for more background on the notion of type.
    ${ }^{5}$ In the context of the a game in Bayesian form, similar to a game in normal form, it is more common to use mixed strategies of type $D\left(S_{i}\right)=D\left(\Theta_{i} \rightarrow A_{i}\right)$, i.e. distributions on pure strategies. However, as any game in Bayesian form can be equivalently modelled as a game in extensive form, and both, behavioural and mixed strategies lead to the same equilibria (under mild assumptions automatically fulfilled in our setup), the difference is immaterial. The reason for our choice of using behavioural strategies is that open games share properties from both standard formulations and it is easier to connect behavioural strategies to our formulation. More on that in Section 4.7.

[^5]:    ${ }^{6}$ One other important reason to use the extensive form is to consider different, typically refined versions of Nash equilibria which take into account the dynamic dimension of games. But this goes beyond our current paper.
    ${ }^{7}$ Note that the original model is not game-theoretic in nature. See [28] Chapter 13 for a discussion and a richer game-theoretic model than presented here.

[^6]:    ${ }^{8}$ See remark in Section 4.7. There are refinements of Bayesian Nash Equilibria which take the dynamics into account. In the context of the lemon market this is relevant if one considers extensions of that model. One prominent example is the role of costly signals that the informed side may use in order to reduce the uncertainty of the buyer. Cf. [28] Chapter 13.C.

[^7]:    ${ }^{9}$ See Chapter 7 in [29] for a general discussion of equilibrium refinements and Chapter 9 in [28] for a discussion of refinements for dynamic games.

