

Compositional Game Theory

Neil Ghani

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

Viktor Winschel

OICOS GmbH
viktor.winschel@googlemail.com

Jules Hedges

University of Oxford
julian.hedges@cs.ox.ac.uk

Philipp Zahn

University of St. Gallen
philipp.zahn@unisg.ch

Abstract

We introduce open games as a compositional foundation of economic game theory. A compositional approach potentially allows methods of game theory and theoretical computer science to be applied to large-scale economic models for which standard economic tools are not practical. An open game represents a game played relative to an arbitrary environment and to this end we introduce the concept of *coutility*, which is the utility generated by an open game and returned to its environment. Open games are the morphisms of a symmetric monoidal category and can therefore be composed by categorical composition into sequential move games and by monoidal products into simultaneous move games. Open games can be represented by string diagrams which provide an intuitive but formal visualisation of the information flows. We show that a variety of games can be faithfully represented as open games in the sense of having the same Nash equilibria and off-equilibrium best responses.

ACM Reference Format:

Neil Ghani, Jules Hedges, Viktor Winschel, and Philipp Zahn. 2018. Compositional Game Theory. In *LICS '18: 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, July 9–12, 2018, Oxford, United Kingdom*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3209108.3209165>

1 Introduction

The concept of *compositionality* is well-known and almost commonplace in computer science, where it is what ultimately allows programmers to scale software to large systems. However, in many other fields compositionality is essentially unknown and hence its benefits are not available. In this paper we introduce compositionality into a field where one might not believe it to be possible: the study of strategic games and Nash equilibria. They are of interest in economics and computer science where optimal decisions are taken by interacting agents with conflicting goals.

In contrast to classical game theory, where games are studied monolithically as one global object, compositional game theory works bottom-up by building large and complex games from smaller components. Such an approach is inherently difficult since the interaction between games has to be considered. Moreover, in the

compositional approach, the equilibria of larger games should be defined from the equilibria of the component games - but *a priori*, there is no reason why this should be possible.

For example, in the prisoner's dilemma game, each player's best option is to defect, although, if they acted as a single agent, they would cooperate. Moreover, if the one-shot prisoner's dilemma game is repeated, then cooperative equilibria become achievable. More generally, the equilibria of a composite game are not necessarily made up from those of the component games, and locally optimal moves are not guaranteed to be globally optimal. In essence, game theory contains *emergent effects* whereby a composite system exhibits behaviours that are not (simple) functions of the behaviours of the components. Accordingly, emergent effects make compositionality very hard to achieve and the existence of a compositional model of game theory is somewhat surprising. In order to arrive at this goal we had to radically reformulate classical game theory from first principles and rebuild it on *open games*.

Open games represent the relationship between different interactions in two dimensions: in sequence if an interaction follows another interaction, and in parallel if interactions take place simultaneously. As such, we follow a path taken in the field of *open systems* [22], and in particular *categorical open systems* [6] where compositional approaches to general systems are studied. Here, systems are modelled as morphisms $f : X \rightarrow Y$ in a symmetric monoidal category, where the objects X and Y describe the *boundaries* of the open system, where it interacts with its environment. This means that systems $f : X \rightarrow Y$ and $f' : X' \rightarrow Y'$ can be composed in parallel using the monoidal product to yield $f \otimes f' : X \otimes X' \rightarrow Y \otimes Y'$, and two systems $f : X \rightarrow Y$ and $g : Y \rightarrow Z$ sharing a common boundary can be composed sequentially by glueing along this boundary to yield $g \circ f : X \rightarrow Z$. Ordinary, closed systems are recovered as scalars [1], i.e. endomorphisms $f : I \rightarrow I$ of the monoidal unit, which represents a trivial boundary. Open games are accordingly the morphisms of a symmetric monoidal category.

A compositional model of game theory does not only have to model a game but also the interactions of the game with all other games and environments. This can be seen as a form of *continuation passing style*. This would still be hardly tractable if the environment of an open game included arbitrary other open games. The crucial technical feature underlying our approach is to describe the behaviour of an open game relative to a simplified notion of an environment which we call a *context*, in which the future is abstracted into a single utility function. In this way, we reduce an arbitrarily complex game to a set of individual decisions. The circularity of a Nash equilibrium, where all players play mutually best replies, is finally handled by the composition operators.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

LICS '18, July 9–12, 2018, Oxford, United Kingdom

© 2018 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-5583-4/18/07...\$15.00

<https://doi.org/10.1145/3209108.3209165>

The theory of open games is based on two main predecessors. Firstly, in [18] games are defined as processes and in [2] the dynamics but not the equilibria are treated compositionally. The second predecessor is the theory of higher order games and selection functions, for example in [4] and [13], which give a theory of equilibria relative to an environment but are not strongly compositional. (Selection functions can however be used to model *goals* of agents compositionally [12].) Combining features of these approaches into a single theory required the innovations mentioned above and led us to discover the idea of an open game. After we developed open games, connections to lenses and the geometry of interaction were noticed respectively by Jeremy Gibbons and Tom Hirschowitz.

Although compositional game theory is partly inspired by game semantics, there is little technical overlap beyond monoidal categories. Game semantics avoids several difficult defining features of game theory by restricting to the 2-player zero-sum setting, which leads to more richly structured (typically *-autonomous) categories. In [11] morphisms between open games are defined to capture *relative* Nash equilibria, directly inspired by relative winning strategies in game semantics, and this leads to deeper connections between open games and game semantics than are apparent from this paper.

We omit proofs in this paper, which can be found in [9] and [11]. We also work over the category of sets to keep notation and overheads to a minimum – for a more categorical account, see once more [9]. The rest of this paper is structured as follows: The next section introduces selection functions as a key ingredient of open games. Section 3 introduces the definition of an open game and discusses its elements, followed by some examples in Section 4. The monoidal category of open games is introduced in Section 5 and the string diagrams attached to this category in Section 6. We finally turn to examples built compositionally: in Section 7 we discuss simultaneous move games and in Section 8 sequential move games. Section 9 concludes the paper with an outlook on further work.

2 Selection functions and higher order games

For reasons of space, we assume the reader knows some basic game theory, such as the definitions of normal-form and extensive-form games and Nash equilibrium. These basic concepts can be found for example in [14] or many online lecture notes. Nevertheless, in this section we introduce enough game theory via selection functions [4, 13] to make the paper self contained.

Definition 2.1. An n -player *higher order simultaneous move game* is defined by the following data:

- For each player $1 \leq i \leq n$, a set X_i of *choices*
- A set R of *outcomes*
- An *outcome function* $q : \prod_{i=1}^n X_i \rightarrow R$
- For each player $1 \leq i \leq n$, a *multi-valued selection function* $\delta_i : (X_i \rightarrow R) \rightarrow \mathcal{P}(X_i)$

In this game, each player simultaneously makes a choice of move $x_i : X_i$. We define a (pure) *strategy* for player i to be just a choice in X_i , and a (pure) *strategy profile* to be a tuple of strategies in $\Sigma := \prod_{i=1}^n X_i$. When all choices are made, the rules of the game determine an outcome $q(x_1, \dots, x_n) : R$. The selection function $\delta_i : (X_i \rightarrow R) \rightarrow \mathcal{P}(X_i)$ defines the set of moves $\delta_i(k)$ that are considered optimal in a *context* $k : X_i \rightarrow R$, which describes the individual decision faced by player i . A context for player i is obtained from the outcome function by fixing strategies for all other

players and is thus the utility function associating, to each potential unilateral deviation x_i of player i , the utility arising from that deviation.

Given a higher order simultaneous move game, we define its *best response relation* $\mathbf{B} \subseteq \Sigma \times \Sigma$ by $(\sigma, \sigma') \in \mathbf{B}$ iff for all players $1 \leq i \leq n$, we have $\sigma'_i \in \delta_i(\lambda(x_i : X_i).q(\sigma[i \mapsto x_i]))$, where $\sigma[i \mapsto x_i]$ is the strategy profile σ apart from the i th player who chooses x_i or, more formally,

$$(\sigma[i \mapsto x_i])_j = \begin{cases} x_i & \text{if } i = j \\ \sigma_j & \text{otherwise} \end{cases}$$

(In game theory this is usually written (σ'_i, σ_{-i}) , but we prefer the more precise notation of computer science.) A *selection equilibrium* is a pure strategy profile σ with $(\sigma, \sigma) \in \mathbf{B}$.

The classical definition of an n -player simultaneous move game in normal form results from this definition as the special case in which the set of outcomes is \mathbb{R}^n and the i th player's selection function is

$$\begin{aligned} \delta_i(k) &= \arg \max(\pi_i \circ k) \\ &= \{x : X_i \mid (k(x))_i \geq (k(x'))_i \text{ for all } x' : X_i\} \end{aligned}$$

In this case the selection equilibria agree with the usual definition of pure strategy Nash equilibrium, and moreover the best response relation is the same. There are many well-known examples of 2-player simultaneous move games with 2 moves each, such as the prisoner's dilemma, matching pennies, battle of the sexes, chicken, etc., defined by different outcome functions $q : \{C, D\}^2 \rightarrow \mathbb{R}^2$. The prisoner's dilemma, for example, is given by the outcome function

$$\begin{aligned} q(C, C) &= (2, 2) & q(C, D) &= (0, 3) \\ q(D, C) &= (3, 0) & q(D, D) &= (1, 1) \end{aligned}$$

and yields the constant best response relation with $(\sigma, (D, D)) \in \mathbf{B}$ for all strategy profiles σ . This happens because D is a *dominant strategy* for both players in the prisoner's dilemma. An extended example can be found in [13] of a higher order simultaneous move game whose selection functions are not of this form, which we will discuss in section 7.

Definition 2.2. An n -player *higher order sequential game* is defined by the same data as a simultaneous move game: sets X_1, \dots, X_n of choices, a set R of outcomes, an outcome function $q : \prod_{i=1}^n X_i \rightarrow R$, and selection functions $\delta_i : (X_i \rightarrow R) \rightarrow \mathcal{P}(X_i)$. A strategy for player i is a function

$$\sigma_i : \prod_{j=1}^{i-1} X_j \rightarrow X_i$$

that chooses a move contingent on the previous moves by other players, and the set of strategy profiles is the set

$$\Sigma = \prod_{i=1}^n \left(\prod_{j=1}^{i-1} X_j \rightarrow X_i \right)$$

of tuples consisting of a strategy for each player. There is an obvious *play function* $\mathbf{P} : \Sigma \rightarrow \prod_{i=1}^n X_i$ producing the sequence of moves resulting from a strategy profile, defined by course-of-values recursion, which in the base case uses $\prod_{j=1}^0 X_j \rightarrow X_1 \cong X_1$. Given a strategy profile σ , we call $\mathbf{P}(\sigma)$ the *strategic play* of σ .

The best response relation $\mathbf{B} \subseteq \Sigma \times \Sigma$ of a higher order sequential game is defined by $(\sigma, \sigma') \in \mathbf{B}$ iff

$$(\mathbf{P}(\sigma[i \mapsto \sigma'_i]))_i \in \delta_i(\lambda(x_i : X_i).q(\mathbf{U}_i(x_i, \sigma)))$$

where the *unilateral deviation operator* \mathbf{U}_i is the evident function $\mathbf{U}_i : X_i \times \Sigma \rightarrow \prod_{j=1}^n X_j$ defined by course-of-values recursion, with $(\mathbf{U}_i(x_i, \sigma))_i = x_i$. Strategy profiles σ with $(\sigma, \sigma) \in \mathbf{B}$ are called *selection equilibria* of the sequential game.

In the case where the selection functions are arg max as defined above, this agrees with the standard definitions of best responses and Nash equilibria. Note that this is a strictly weaker definition than that of *optimal strategies* from [4], which specialises to *subgame perfect equilibrium* (a strengthening of Nash equilibrium) when the selection functions are arg max [5].


3 Open games

The primary objects of study in compositional game theory are called *open games*. We start by giving the definition, and in the remainder of this section we discuss its interpretation.

Definition 3.1. Let X, S, Y, R be sets. An *open game* $\mathcal{G} : (X, S) \rightarrow (Y, R)$ is defined to be a 4-tuple $\mathcal{G} = (\Sigma_{\mathcal{G}}, \mathbf{P}_{\mathcal{G}}, \mathbf{C}_{\mathcal{G}}, \mathbf{B}_{\mathcal{G}})$, where

- $\Sigma_{\mathcal{G}}$ is a set, called the set of *strategy profiles* of \mathcal{G}
- $\mathbf{P}_{\mathcal{G}} : \Sigma_{\mathcal{G}} \times X \rightarrow Y$ is called the *play function* of \mathcal{G}
- $\mathbf{C}_{\mathcal{G}} : \Sigma_{\mathcal{G}} \times X \times R \rightarrow S$ is called the *coplay function* of \mathcal{G}
- $\mathbf{B}_{\mathcal{G}} : X \times (Y \rightarrow R) \rightarrow \text{Rel}(\Sigma_{\mathcal{G}})$ is called the *best response function* of \mathcal{G}

$\text{Rel}(\Sigma_{\mathcal{G}})$ is the meet-semilattice of all endo-relations $R \subseteq \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{G}}$. In general, we impose no conditions whatsoever on these components. In practice, however, we are most interested in those open games which are generated by certain constructions, some of which are defined in this paper. There are class-many open games of a fixed type (which causes the category of open games to be locally large), but only set-many after restricting to those generated by a set of constructions. (We could also simply bound the cardinality of $\Sigma_{\mathcal{G}}$.) We will represent a general open game $\mathcal{G} : (X, S) \rightarrow (Y, R)$ in diagrammatic form (where time flows from left to right) as



In Section 5 we will build on this and discuss the graphical syntax of open games.

We interpret X as the type of *observations* that can be made by \mathcal{G} that inform the choice of a strategy and hence action, and Y as the type of *moves* or *choices*. By that a game \mathcal{G} is a *process* that maps observations X to choices Y . The types R and S , on the other hand, are ‘dual’ or ‘contravariant’ types and this is indicated above by the arrows flowing in the reverse direction. We think of R as the type of *utility* (type of values about which the players in \mathcal{G} have preferences) that actions might generate. Thus utility functions that arise in standard game theory are simply functions $Y \rightarrow R$. The type S is dually called the *coutility* since it represents values that are returned to the calling environment by the game so that they can become utility for other processes. This is seen clearly and formally within the definition of composition of open games where the utility of one game (acting as the environment for the other) is

computed via the coplay function of the other game. We point this out explicitly when discussing the composition of games.

The most straightforward parts of the definition of open games are the first two components. It is intuitive that a game has a set of strategy profiles and that, given a strategy profile and an observation, we can run the strategy profile on the observation to obtain a choice. To give a simple concrete example, suppose $Y = A \times B$, and define

$$\Sigma = (X \rightarrow A) \times (X \times A \rightarrow B)$$

and $\mathbf{P}((\sigma_1, \sigma_2), x) = (\sigma_1(x), \sigma_2(x, \sigma_1(x)))$. This represents a two-player game of *perfect information*: the value x is the input which the first player observes and then chooses a using the strategy σ_1 . The second player observes both x and a and chooses b using σ_2 . The pair (a, b) is taken to be the output.

To gain an intuition for coplay, let us first consider a very simple situation. You receive your monthly income and upon observing your bonus (your X), you decide to buy a bottle of champagne. You go to a wine shop and for the given price, you buy a bottle (your Y), which gives you a certain utility (your R). Here, the connection between choice and reaction, between what you do and what comes back to you, is rather immediate. Deciding is simple.

Open games model situations where the connection between your action and what comes back to you is left open - in the same way as with selection functions: by allowing for *all* possible contexts. The only prerequisite for contexts is being well-typed. Coming back to the example, suppose, instead of buying the bottle yourself, you give an amount of money to your friend, a champagne aficionado, with the request to buy a bottle for you. So, your action gives her some money, and you expect a bottle of champagne and by that some utility back from her. She observes your bonus and the amount of money given to her (your Y , her X), and will take some action (her Y) that will bring back a bottle of champagne from a yet unspecified environment (her R). Ultimately, she will hand back a bottle of champagne (her S , your R) which will create utility for you. This handing back of utility is computed through her coplay function which computes the coutility to hand back.

Note, in the example we used the concept of *utility*, which by definition is the real number maximised by classical game-theoretic agents. There is no requirement in the definition of an open game that outcomes are either real numbers, or a linear resource such as money nor that some utility is maximised at all. Indeed, slightly changing the example, we could have defined coutility as belonging to a certain set of possible types of wine. This flexibility (which is inherited from the theory of selection functions) may be seen as a necessary side-effect of obtaining compositionality, but can itself be useful in modelling [12, 13].

Possibly the most important part of the definition is the best response relation, which is defined relative to an arbitrary *context* - as in the case of games modelled via selection functions. Hence, an open game also has Nash equilibria relative to an arbitrary context. A context consists of a *state*, which says what happened in the past, and a utility function, expressed as a *continuation*, which says what will happen in the future. Compositional game theory relies on the observation that such *relative* best response relations provide a strategic representation of games that can be composed.

Relationship with Lenses: Pairs of functions $X \rightarrow Y$ and $X \times R \rightarrow S$ are equivalent to *polymorphic lenses* $(X, S) \rightarrow (Y, R)$ [10, 19].

Moreover, all open games can be built from lenses in a way we now describe. Observe that

- An element $x : X$ is given by a lens $(1, 1) \rightarrow (X, S)$, where $1 = \{*\}$. Call such lenses *states* $\text{St}(X, S)$
- A function $Y \rightarrow R$ is a lens $(Y, R) \rightarrow (1, 1)$. Call such lenses *costates* $\text{CoSt}(Y, R)$

Lemma 3.2. *An open game $\mathcal{G} : (X, S) \rightarrow (Y, R)$ is exactly*

- *A family of lenses, that is a set $\Sigma_{\mathcal{G}}$ and, for each $\sigma : \Sigma_{\mathcal{G}}$, a lens $\mathcal{G}_{\sigma} : (X, S) \rightarrow (Y, R)$.*
- *$\mathbf{B}_{\mathcal{G}} : \text{St}(X, S) \times \text{CoSt}(Y, R) \rightarrow \text{Rel}(\Sigma_{\mathcal{G}})$*

In such a situation, we often save notational overhead by writing the lens \mathcal{G}_{σ} as σ , as in $(X, S) \xrightarrow{\sigma} (Y, R)$.

Lenses play a key role in the development of open games as they hide many details which otherwise severely pollute the presentation. More concretely, they encapsulate the purely algebraic parts of open games, leaving us to focus on the strategic behaviour of the best-response function. They also ensure all reasoning about open games can take place diagrammatically in the category of lenses. As a result, more recent work on open games has heavily exploited this connection [11].

Relationship with Geometry of Interaction: It has been pointed out to us that there is a connection between the geometry of interaction (GoI) and open games. A central construction within GoI, the Int-construction, takes a traced monoidal category \mathbf{C} and constructs another category $I(\mathbf{C})$ whose objects are pairs of objects of \mathbf{C} , and whose morphisms $(X, S) \rightarrow (Y, R)$ are maps $X \otimes R \rightarrow Y \otimes S$. If \mathbf{C} is cartesian, this is equivalent to two functions $X \times R \rightarrow Y$ and $X \times R \rightarrow S$. If the former function does not use the input R , we get precisely a lens. This restriction means that open games are not completely symmetric and this explains why a trace operator is not needed to define composition — one can simply calculate the forwards/covariant part of the composition and use that to calculate the backwards/contravariant part.

When the base category is traced cartesian monoidal, both lenses and the Int-construction are defined, and the two composition operators agree. This means that there is an identity-on-objects functor from the category of lenses to $I(\mathbf{C})$, and this suggests it will be possible to build a more symmetrical theory of open games with *computable* strategies.

4 Examples of open games

An open game $\mathcal{G} : (1, 1) \rightarrow (1, 1)$, where $1 = \{*\}$, consists (up to isomorphism of sets) of a set $\Sigma_{\mathcal{G}}$ of strategy profiles and a best response relation $\mathbf{B}_{\mathcal{G}} \subseteq \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{G}}$. Since $(1, 1)$ will turn out to be the monoidal unit of the monoidal category of open games, we call \mathcal{G} a *scalar open game*, following the terminology of [1].

Any existing notion of ‘game’ from which we can define a best response function can be encoded as a scalar open game. For example, consider the prisoner’s dilemma from Section 2. This can be represented as a scalar open game $\mathcal{G} : (1, 1) \rightarrow (1, 1)$ with $\Sigma_{\mathcal{G}} = \{C, D\}^2$ with the constant best response relation $\mathbf{B}_{\mathcal{G}}$ defined by $(\sigma, \sigma') \in \mathbf{B}_{\mathcal{G}}$ iff $\sigma' = (D, D)$. (The question of how much information about a normal-form or extensive-form game is preserved by the representation as a scalar open game is a subtle one, which we can unfortunately not discuss here for space reasons.)

In this paper we will show that open games can be used to build such scalars compositionally. In the remainder of this section we

will define families of open games that we consider *atomic*, in the sense that they are not built compositionally from smaller games. These families are *decisions*, *functions*, and *countits*.

A decision is an open game that represents a single choice made by an agent.

Definition 4.1. Let X and Y be sets. A (utility-maximising) *decision* $\mathcal{D} : (X, 1) \rightarrow (Y, \mathbb{R})$ is an open game defined by the following data:

- $\Sigma_{\mathcal{D}} = X \rightarrow Y$
- $\mathbf{P}_{\mathcal{D}}(\sigma, x) = \sigma(x)$
- $\mathbf{C}_{\mathcal{D}}(\sigma, x, r) = *$
- $(\sigma, \sigma') \in \mathbf{B}_{\mathcal{D}}(x, k)$ iff $\sigma'(x) \in \arg \max k$

A (pure) strategy for a single decision is a function that maps possible observations that can be made by the agent, to possible choices. Such a strategy is considered optimal in the context (x, k) iff it maps the current state x to a maximising point of the current continuation k .

The reason that the pre-deviation strategy σ plays no role in $(\sigma, \sigma') \in \mathbf{B}_{\mathcal{D}}(x, k)$ is that σ is considered the strategy played by all *other* players besides the one currently under consideration, and so it plays no role in a 1-player open game such as \mathcal{D} . Put it in another way: in a one-player game there is nothing for σ' to be a best response *to*.

The definition of a decision assumes an agent who maximises real-valued utility. However, there is nothing inherent in the definition of an open game that says this must be the case. Indeed, one can argue that the additional generality is necessary to be compositional: An aggregate of two maximising agents, if modelled as a selection function-like object, is not necessarily maximising. The prisoner’s dilemma is a standard example of an aggregate with behaviour that is not (globally) maximising, that is to say, if the two players were modelled as a single entity, they would choose (C, C) . This is the sense in which selection functions are a theoretical precursor to open games.

Definition 4.2. Let X, Y and R be sets, and let $\delta : (Y \rightarrow R) \rightarrow \mathcal{P}(Y)$ be a multi-valued selection function. We define an open game $\mathcal{D}_{\delta} : (X, 1) \rightarrow (Y, R)$ by the following data:

- $\Sigma_{\mathcal{D}_{\delta}} = X \rightarrow Y$
- $\mathbf{P}_{\mathcal{D}_{\delta}}(\sigma, x) = \sigma(x)$
- $\mathbf{C}_{\mathcal{D}_{\delta}}(\sigma, x, r) = *$
- $(\sigma, \sigma') \in \mathbf{B}_{\mathcal{D}_{\delta}}(x, k)$ iff $\sigma'(x) \in \delta(k)$

We will give an example of an open game with non-utility-maximising decisions in Section 7.

Decisions are the only atomic open games that play a role in strategic reasoning. We formalise this with the following definition.

Definition 4.3. Let $\mathcal{G} : (X, S) \rightarrow (Y, R)$ be an open game. We call \mathcal{G} *strategically trivial* if it satisfies the following two conditions:

- $|\Sigma_{\mathcal{G}}| = 1$ (say, $\Sigma_{\mathcal{G}} = \{*\}$)
- $(*, *) \in \mathbf{B}_{\mathcal{G}}(x, k)$ for all contexts (x, k)

The first condition says that there is exactly one strategy, and so there is no choice to be made. The second condition says that this trivial strategy can never fail to be in equilibrium. The idea behind this is that if a strategy profile fails to be an equilibrium, it should always be because some player has an incentive to deviate. Strategically trivial open games could also be called zero-player open games.

Definition 4.4. Let $f : X \rightarrow Y$ and $g : R \rightarrow S$ be functions. We define a strategically trivial open game $(f, g) : (X, S) \rightarrow (Y, R)$ with play function $\mathbf{P}_{(f,g)}(x) = f(x)$ and coplay function $\mathbf{C}_{(f,g)}(*, x, r) = g(r)$.

As a special case of this, a function $f : X \rightarrow Y$ can be ‘lifted’ to an open game in two ways: covariantly as $(f, \text{id}_1) : (X, 1) \rightarrow (Y, 1)$, or contravariantly as $(\text{id}_1, f) : (1, Y) \rightarrow (1, X)$.

The final class of atomic open games are the *counits*.

Definition 4.5. Let X be a set. We define a strategically trivial open game $\varepsilon_X : (X, X) \rightarrow (1, 1)$ called a *cunit*, with play function $\mathbf{P}_{\varepsilon_X}(*, x) = *$ and coplay function $\mathbf{C}_{\varepsilon_X}(*, x, *) = x$.

Recall that backward-flowing values are ‘teleological’, that is, they are future values about which agents are reasoning. The role of counits is to identify an ordinary forward-flowing value as the value *about which* some past agent is reasoning. This plays an important role in the diagrammatic language of open games in the next section. Note there is no dual unit of type $(1, 1) \rightarrow (X, X)$ (the reader might like to try to define one) which is a reflection of the fact that game theory is not symmetric in its forward and backward looking facets. Mathematically, open games will not form a compact closed category.

5 The monoidal category of open games

In this section we will define a pair of operators for composing open games, categorical composition and monoidal product, which correspond to *sequential play* and *simultaneous play*. These two operators make open games into the morphisms of a symmetric monoidal category (after quotienting by isomorphisms of strategy sets). As morphisms of a monoidal category open games can also be denoted by string diagrams, and we introduce this diagrammatic language in the next section. While these two operators form the core, we do not claim that they are a *complete* basis of operators for building open games in any formal or informal sense. Indeed, other operators are discussed in [8].

First we give the definition of categorical composition. This is a form of sequential composition in which the choice made by the first component is *hidden*, visible only to the second component but not to the outside. Sequential play is more intuitive when the choices made by both components are visible; this will be recovered as a derived operator in Section 8.

Definition 5.1. Given a pair of open games $\mathcal{G} : (X, S) \rightarrow (Y, R)$ and $\mathcal{H} : (Y, R) \rightarrow (Z, Q)$, we define their composition $\mathcal{H} \circ \mathcal{G} : (X, S) \rightarrow (Z, Q)$ as follows. The set of strategy profiles is the cartesian product $\Sigma_{\mathcal{H} \circ \mathcal{G}} = \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}$. The play function composes simply by composition of functions:

$$\mathbf{P}_{\mathcal{H} \circ \mathcal{G}}((\sigma, \tau), x) = \mathbf{P}_{\mathcal{H}}(\tau, \mathbf{P}_{\mathcal{G}}(\sigma, x))$$

The coplay function composes as follows:

$$\mathbf{C}_{\mathcal{H} \circ \mathcal{G}}((\sigma, \tau), x, q) = \mathbf{C}_{\mathcal{G}}(\sigma, x, \mathbf{C}_{\mathcal{H}}(\tau, \mathbf{P}_{\mathcal{G}}(\sigma, x), q))$$

The best response relation $((\sigma, \tau), (\sigma', \tau')) \in \mathbf{B}_{\mathcal{H} \circ \mathcal{G}}(x, k)$ holds iff $(\sigma, \sigma') \in \mathbf{B}_{\mathcal{G}}(x, k')$ and $(\tau, \tau') \in \mathbf{B}_{\mathcal{H}}(\mathbf{P}_{\mathcal{G}}(\sigma, x), k)$, where $k' : Y \rightarrow R$ is defined by $k'(y) = \mathbf{C}_{\mathcal{H}}(\tau, y, k(\mathbf{P}_{\mathcal{H}}(\tau, y)))$.

Since the set of strategy profiles is usually a tuple consisting of a strategy for each decision, the condition $\Sigma_{\mathcal{H} \circ \mathcal{G}} = \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}$ corresponds roughly to saying that the set of decisions in $\mathcal{H} \circ \mathcal{G}$ is the disjoint union of the decisions in \mathcal{G} and \mathcal{H} . The play function

says that to play the sequential composition $\mathcal{H} \circ \mathcal{G}$ in state x with strategy profile (σ, τ) , we first play \mathcal{G} with σ in state x , obtaining a state y for \mathcal{H} , which we then play with τ .

The formula for composing coplay functions is hard to understand intuitively, but successfully captures the informal intuition given for coplay in section 3. Alternatively, it can be seen as composition of lenses [10].

Finally, we give the conditions for a strategy profile (σ', τ') to be a best response to (σ, τ) in a context (x, k) . This means that for each player in \mathcal{G} , σ' must be rational assuming that all other players in \mathcal{G} play σ and all players in \mathcal{H} play τ , and also that for each player in \mathcal{H} , τ' must be rational assuming that all players in \mathcal{G} play σ and all other players in \mathcal{H} play τ . (This is Nash equilibrium-like reasoning.) We can apply the compositionally-known best response relations for \mathcal{G} and \mathcal{H} , after using these assumptions to appropriately modify the context. For \mathcal{H} the continuation remains k , and the state x is modified to $\mathbf{P}_{\mathcal{G}}(\sigma, x)$ using the assumption that players in \mathcal{G} play σ . For \mathcal{G} the state remains x , and the continuation is modified to k' , using the interpretation of the coplay function $\mathbf{C}_{\mathcal{H}}$ as the utility passed backward from \mathcal{H} to \mathcal{G} , with the assumption that players in \mathcal{H} play τ .

Open games trivially cannot form a category, because this composition operator is not associative on the nose: the strategy profiles are

$$\begin{aligned} \Sigma_{\mathcal{I} \circ (\mathcal{H} \circ \mathcal{G})} &= (\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}) \times \Sigma_{\mathcal{I}} \\ &\neq \Sigma_{\mathcal{G}} \times (\Sigma_{\mathcal{H}} \times \Sigma_{\mathcal{I}}) = \Sigma_{(\mathcal{I} \circ \mathcal{H}) \circ \mathcal{G}} \end{aligned}$$

There are three approaches to this problem. The first, which is perfectly successful in practice, is to simply ignore it and informally work up to isomorphic strategy sets. The second, which is attractive from a theoretical point of view, is to define a bicategory of open games in which the 2-cells are functions between strategy sets that suitably commute with the remaining structure. The reason we do not take this approach is that monoidal bicategories are notoriously complicated, and generalising to a monoidal double category (whose axioms are typically much easier to verify [21]) would take us too far afield. (This is carried out in [11].) Therefore in this paper we take the third approach, which is to define a suitable equivalence relation on open games and then a category whose morphisms are equivalence classes. This is equivalent to first defining a bicategory, and then obtaining a 1-category by quotienting by invertible 2-cells.

Definition 5.2. Let $\mathcal{G}_1, \mathcal{G}_2 : (X, S) \rightarrow (Y, R)$ be open games. We write $\mathcal{G}_1 \sim \mathcal{G}_2$ if there is a bijection $i : \Sigma_{\mathcal{G}_1} \rightarrow \Sigma_{\mathcal{G}_2}$ such that

- $\mathbf{P}_{\mathcal{G}_1}(\sigma, x) = \mathbf{P}_{\mathcal{G}_2}(i(\sigma), x)$ for all $x : X$ and $\sigma : \Sigma_{\mathcal{G}_1}$
- $\mathbf{C}_{\mathcal{G}_1}(\sigma, x, r) = \mathbf{C}_{\mathcal{G}_2}(i(\sigma), x, r)$ for all $x : X, r : R$ and $\sigma : \Sigma_{\mathcal{G}_1}$
- $(\sigma, \sigma') \in \mathbf{B}_{\mathcal{G}_1}(x, k)$ iff $(i(\sigma), i(\sigma')) \in \mathbf{B}_{\mathcal{G}_2}(x, k)$ for all $x : X, k : Y \rightarrow R$ and $\sigma, \sigma' : \Sigma_{\mathcal{G}_1}$

Proposition 5.3. For each type $(X, S) \rightarrow (Y, R)$, \sim is an equivalence relation on the class of open games of that type.

The quotient under \sim identifies open games with isomorphic strategy profiles and best responses. This is close in spirit to the concept of *best response equivalence* of games in classical game theory [16, p. 52f], but has some strange consequences. For example, let \mathcal{G} be the scalar open game representing the prisoner’s dilemma from the previous section, with $\Sigma_{\mathcal{G}} = \{C, D\}^2$ and $(\sigma, \sigma') \in \mathbf{B}_{\mathcal{G}}$ iff $\sigma' = (D, D)$. Now consider a 1-player game with 4 choices

$A = \{1, 2, 3, 4\}$ and utility function $k : A \rightarrow \mathbb{R}$ given by $k(x) = x$. Encoding this as a scalar open game \mathcal{H} yields $\Sigma_{\mathcal{H}} = A$ and $(\sigma, \sigma') \in \mathbf{B}_{\mathcal{H}}$ iff $x' = 4$, since 4 is again a dominant strategy for the player. Then $\mathcal{G} \sim \mathcal{H}$, and so in the quotient they will be *equal* elements of the monoid of scalars, despite the fact that they represent games with different numbers of players. However, this is a technical consequence of working with an equivalence relation, and does not happen if we instead use a bicategory or double category of open games.

Proposition 5.4. \circ is well-defined on equivalence classes, that is to say, if $\mathcal{G} \sim \mathcal{G}'$ and $\mathcal{H} \sim \mathcal{H}'$ then $\mathcal{H} \circ \mathcal{G} \sim \mathcal{H}' \circ \mathcal{G}'$.

Definition 5.5. For each object (X, S) , the identity open game $\text{id}_{(X,S)} : (X, S) \rightarrow (X, S)$ is the function $(\text{id}_X, \text{id}_S)$, that is, the strategically trivial open game with $\mathbf{P}_{\text{id}_{(X,S)}}(*, x) = x$ and $\mathbf{C}_{\text{id}_{(X,S)}}(*, x, s) = s$.

Proposition 5.6. There is a category **Game** whose objects are pairs of sets and whose morphisms are equivalence classes of open games. The composition is \circ and the identity on (X, S) is the equivalence class of $\text{id}_{(X,S)}$.

In the previous section we defined a strategically trivial open game $(f, g) : (X, S) \rightarrow (Y, R)$ given functions $f : X \rightarrow Y$ and $g : R \rightarrow S$. This defines a functor $(-, -) : \mathbf{Set} \times \mathbf{Set}^{\text{op}} \rightarrow \mathbf{Game}$.

Next we define the monoidal product of open games, which corresponds to simultaneous play.

Definition 5.7. Let $\mathcal{G}_1 : (X_1, S_1) \rightarrow (Y_1, R_1)$ and $\mathcal{G}_2 : (X_2, S_2) \rightarrow (Y_2, R_2)$ be open games. We define an open game $\mathcal{G}_1 \otimes \mathcal{G}_2 : (X_1 \times X_2, S_1 \times S_2) \rightarrow (Y_1 \times Y_2, R_1 \times R_2)$ as follows:

- The set of strategy profiles is $\Sigma_{\mathcal{G}_1 \otimes \mathcal{G}_2} = \Sigma_{\mathcal{G}_1} \times \Sigma_{\mathcal{G}_2}$
- The play function is

$$\mathbf{P}_{\mathcal{G}_1 \otimes \mathcal{G}_2}((\sigma_1, \sigma_2), (x_1, x_2)) = (\mathbf{P}_{\mathcal{G}_1}(\sigma_1, x_1), \mathbf{P}_{\mathcal{G}_2}(\sigma_2, x_2))$$

- The coplay function is

$$\begin{aligned} \mathbf{C}_{\mathcal{G}_1 \otimes \mathcal{G}_2}((\sigma_1, \sigma_2), (x_1, x_2), (r_1, r_2)) \\ = (\mathbf{C}_{\mathcal{G}_1}(\sigma_1, x_1, r_1), \mathbf{C}_{\mathcal{G}_2}(\sigma_2, x_2, r_2)) \end{aligned}$$

- The relation

$$((\sigma_1, \sigma_2), (\sigma'_1, \sigma'_2)) \in \mathbf{B}_{\mathcal{G}_1 \otimes \mathcal{G}_2}((x_1, x_2), k)$$

holds iff the relations $(\sigma_1, \sigma'_1) \in \mathbf{B}_{\mathcal{G}_1}(x_1, k_1)$ and $(\sigma_2, \sigma'_2) \in \mathbf{B}_{\mathcal{G}_2}(x_2, k_2)$ both hold, where $k_1 : Y_1 \rightarrow R_1$ and $k_2 : Y_2 \rightarrow R_2$ are defined by

$$k_1(y_1) = \pi_1(k(y_1, \mathbf{P}_{\mathcal{G}_2}(\sigma_2, x_2)))$$

$$k_2(y_2) = \pi_2(k(\mathbf{P}_{\mathcal{G}_1}(\sigma_1, x_1), y_2))$$

Lemma 5.8. \otimes is well-defined on equivalence classes.

Lemma 5.9. \otimes defines a bifunctor $\mathbf{Game} \times \mathbf{Game} \rightarrow \mathbf{Game}$.

Theorem 5.10. **Game** is a monoidal category in which the monoidal product on objects is $(X_1, S_1) \otimes (X_2, S_2) = (X_1 \times X_2, S_1 \times S_2)$ and that on morphisms is previously defined. The monoidal unit is the object $I = (1, 1)$.

The structure morphisms of the monoidal category are inherited from the monoidal category $\mathbf{Set} \times \mathbf{Set}^{\text{op}}$ via the functor $(-, -)$, where **Set** is cartesian monoidal. **Game** is moreover symmetric monoidal, with the symmetry inherited from $\mathbf{Set} \times \mathbf{Set}^{\text{op}}$.

A Lens-theoretic View: A cleaner approach arises if one factors the definition of parallel and sequential composition via the use of lenses. First note that the above definitions restrict to lenses meaning that the category of lenses is symmetric monoidal. Just as open games can be defined in terms of lenses, the categorical structure of open games can similarly be defined in terms of the simpler categorical structure of lenses.

Lemma 5.11. Let $\mathcal{G} : (X, S) \rightarrow (Y, R)$ and $\mathcal{H} : (Y, R) \rightarrow (Z, T)$ be open games. Then the composite $\mathcal{H} \circ \mathcal{G}$ is the family of lenses indexed by $\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}$ with the pair (σ, τ) indexing the lens

$$(X, S) \xrightarrow{\sigma} (Y, R) \xrightarrow{\tau} (Z, T)$$

Given a state $h : (1, 1) \rightarrow (X, S)$ and a costate $k : (Z, T) \rightarrow (1, 1)$, then $((\sigma, \tau), (\sigma', \tau')) \in \mathbf{B}_{\mathcal{H} \circ \mathcal{G}}(x, k)$ holds iff $(\sigma, \sigma') \in \mathbf{B}_{\mathcal{G}}(h, k \circ \tau)$ and $(\tau, \tau') \in \mathbf{B}_{\mathcal{H}}(\sigma \circ h, k)$.

Notice how lens composition hides all the details of how play and coplay functions knit together in the composite game. In particular, in the definition of composition, the function


$$k'(y) = \mathbf{C}_{\mathcal{H}}(\tau, y, k(\mathbf{P}_{\mathcal{H}}(\tau, y)))$$

is merely the lens composite $k \circ \tau$. A similarly simplified construction of the monoidal product of open games via the monoidal product of lenses can also be given. Observing this vastly simplifies the proof that open games form a symmetric monoidal category; compare the direct proof in [9] and the proof using lenses in [11].

6 String diagrams


We will now informally present the string diagram language for open games. A formal presentation can be found in [10], which proves a coherence theorem by which we can *define* an open game by its string diagram, given interpretations of the atomic open games. We refer the reader to [20] for a summary of graphical languages of this sort. The language of open games is an extension of those of symmetric monoidal categories, and similar to (but not exactly) a fragment of compact closed categories.

The key idea of the string diagram language is to treat the object (X, S) as a formal tensor product $X \otimes S^*$, where $-^*$ is a duality that is defined on objects but not on arbitrary open games. Diagrammatically we represent this duality by an orientation on strings, and thus a general open game $\mathcal{G} : (X, S) \rightarrow (Y, R)$ is denoted by




More formally, we allow individual strings to represent *covariant objects* of the form $(X, 1)$, and *contravariant objects* of the form $(1, S)$. Then, up to isomorphism, a general object can be written as a tensor product $(X, S) \cong (X, 1) \otimes (1, S)$ of a covariant and a contravariant object. If we define a duality operation on arbitrary objects by $(X, S)^* = (S, X)$, then $(X, S) \cong (X, 1) \otimes (S, 1)^*$. This justifies the informal statement $(X, S) = X \otimes S^*$, because we can identify **Set** with a symmetric monoidal subcategory $\mathbf{Set} \hookrightarrow \mathbf{Game}$ by identifying X with $(X, 1)$ and f with (f, id_1) .

Notice that since we also have $(X, S) \cong (S, 1)^* \otimes (X, 1)$, the open game \mathcal{G} can equally be denoted




That is, the relative ordering of covariant and contravariant parts of an object does not matter. More formally, the objects $X \otimes S^*$ and $S^* \otimes X$ are equal in the strictification of **Game**, and the symmetry s_{X, S^*} is an identity.

Corresponding to each atomic open game we have a corresponding ‘atomic’ string diagram, which we compose by the usual operations of end-to-end and side-by-side juxtaposition. For example, a utility-maximising decision $\mathcal{D} : (X, 1) \rightarrow (Y, \mathbb{R})$ corresponds to a node




If X is a 1-element set, we further restrict this to




with the usual (purely syntactic) convention of using a triangle for morphisms into or out of the monoidal unit.

Given a function $f : X \rightarrow Y$, its covariant lifting $(f, \text{id}_1) : (X, 1) \rightarrow (Y, 1)$ and its contravariant lifting $(\text{id}_1, f) : (1, Y) \rightarrow (1, X)$ are respectively denoted




This is analogous to the use of *rotation* for the transpose of a linear map used by Bob Coecke and others, for example in [3]. While transposition in a compact closed category is tensor-reversing, i.e. $(X \otimes Y)^* \cong Y^* \otimes X^*$, for us it is more convenient (ultimately because our objects are pairs of sets) to have $(X \otimes Y)^* = X^* \otimes Y^*$, leading to a notation with reflection rather than rotation.

The deleting function $X \rightarrow 1$ and diagonal function $X \rightarrow X^2$ lift to give a cocommutative comonoid on every covariant object, and a commutative monoid on every contravariant object. We give these the special syntax




The final atomic open games that we must give representations to are the counits $\epsilon_X : (X, X) \rightarrow I$. This is denoted by a bending wire




Since there is no natural strategically trivial open game $I \rightarrow (X, X)$, we do not allow wires to bend in the opposite direction in our string diagrams.

Covariant functions, contravariant functions and counits are related by the *counit law*, stating that the string diagrams



denote equal open games. That is to say, the diagram of open games



commutes (where the monoidal structure morphisms are implicit).

Given a diagram built from the pieces we have described, which does not contain any wire bending in the illegal direction, we can compositionally build an open game where

- Decision and function nodes are interpreted as the corresponding atomic open game
- Side-by-side and end-to-end composition of diagrams is interpreted as monoidal product and categorical composition of open games
- A backwards-bending wire is interpreted as the corresponding counit

The coherence theorem for teleological categories [10] states that the resulting open game is invariant under topological manipulations of the diagram, including rotating function nodes around a bend using the counit law, provided that the new diagram does not contain a wire bending in the illegal direction. Several examples of interpreting a diagram as an open game can be seen in the next two sections.

7 Simultaneous move games

In this section and the next we will apply the theory of the previous sections to demonstrate that various classes of games can be represented as open games and can be built compositionally.

We begin with simultaneous move games. The decision $\mathcal{D}_{1, X_i} : I \rightarrow (X_i, \mathbb{R})$ represents an agent who makes a choice from a set X_i in order to maximise a real number.

Theorem 7.1. *Let*

$$\mathcal{G} := \bigotimes_{i=1}^n \mathcal{D}_{1, X_i} : I \rightarrow \left(\prod_{i=1}^n X_i, \mathbb{R}^n \right)$$

be a monoidal product of decisions. Then the set of strategy profiles of \mathcal{G} is equal to the set of pure strategy profiles of a simultaneous move game with sets of choices X_i , namely $\Sigma_{\mathcal{G}} = \prod_{i=1}^n X_i$ and, for any function $q : \prod_{i=1}^n X_i \rightarrow \mathbb{R}^n$, the relation $\mathbf{B}_{\mathcal{G}}(, q) \subseteq \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{G}}$ is precisely the best response relation for the simultaneous move game*




Figure 1. String diagram for simultaneous move game

with outcome function q (and, hence, the fixpoints of $\mathbf{B}_{\mathcal{G}}(*, q)$ are the pure strategy Nash equilibria).

In particular, by the associativity of the monoidal product, the monoidal product of an m -player and an n -player open game is an $m + n$ -player open game.

Given a particular utility function $q : \prod_{i=1}^n X_i \rightarrow \mathbb{R}^n$, consider the covariant function

$$(q, 1) : \left(\prod_{i=1}^n X_i, 1 \right) \rightarrow (\mathbb{R}^n, 1)$$

We take the monoidal product of this with an identity morphism and then post-compose with a counit, to yield the strategically trivial open game

$$\left(\prod_{i=1}^n X_i, \mathbb{R}^n \right) \xrightarrow{(q, 1) \otimes \text{id}_{(1, \mathbb{R}^n)}} (\mathbb{R}^n, \mathbb{R}^n) \xrightarrow{\varepsilon_{\mathbb{R}^n}} (1, 1)$$

By the counit law, this can be equivalently written as

$$\varepsilon_{\prod_{i=1}^n X_i} \circ (\text{id}_{(\prod_{i=1}^n X_i, 1)} \otimes (1, q))$$

Now consider the scalar open game

$$(1, 1) \xrightarrow{\bigotimes_{i=1}^n \mathcal{D}_{1, X_i}} \left(\prod_{i=1}^n X_i, \mathbb{R}^n \right) \xrightarrow{(q, 1) \otimes \text{id}_{(1, \mathbb{R}^n)}} (\mathbb{R}^n, \mathbb{R}^n) \xrightarrow{\varepsilon_{\mathbb{R}^n}} (1, 1)$$

This scalar open game has the property that its best response relation $\mathbf{B}(*, *)$, for the unique context $(*, *)$, is precisely the best response function for the simultaneous move game with outcome function q .

For small values of n , we can visualise this scalar open game as a string diagram. For example, when $n = 2$ the corresponding string diagram is depicted in Figure 1. Here, for the first time we can see how the information flow in a game is visualised with a string diagram: the utility generated by the utility function is ‘fed back’ to each agent via a counit.

The previous results can be strengthened to an arbitrary higher order game (introduced in Section 2) using the open games \mathcal{D}_{δ} associated to a selection function from Section 4, in which case the fixpoints of the best response relation are selection equilibria. An interesting special case of this is depicted in Figure 2, in which




Figure 2. Simultaneous move game with non-utility-maximising players

the outcome that is ‘optimised’ by each player is nothing but the choice of the other player. (Ignoring the types, this string diagram arises from Figure 1 by replacing q with a symmetry, i.e. crossing wires; this provides a nontrivial example of reasoning about the equivalence of games from topological manipulations of string diagrams.) Without imposing an order relation on the set of choices, it is not possible to interpret \mathcal{D}_1 and \mathcal{D}_2 as utility-maximising decisions. Instead, by interpreting \mathcal{D}_1 and \mathcal{D}_2 in suitable ways we can obtain directly analogous results to the Keynes beauty contest example in [13].

We will consider three different ways of interpreting \mathcal{D}_1 and \mathcal{D}_2 , which result in three different games. If $\mathcal{D}_1 = \mathcal{D}_2$ are both \mathcal{D}_{fix} , where $\text{fix} : (X \rightarrow X) \rightarrow \mathcal{P}(X)$ is the selection function $\text{fix}(k) = \{x : X \mid x = k(x)\}$, the resulting scalar open game is a coordination game. In particular, if $X = \{A, B\}$ then the best response relation is the same as *Meeting in New York*, the 2-player simultaneous move game with utility maximising players and outcome function

$$q(x, y) = \begin{cases} (1, 1) & \text{if } x = y \\ (0, 0) & \text{if } x \neq y \end{cases}$$

In particular, the pure Nash equilibria and the fixpoints of $\mathbf{B}_{\mathcal{G}}(*, *)$ are (A, A) and (B, B) .

Next, we interpret both \mathcal{D}_1 and \mathcal{D}_2 as the open game lifted from the anti-fixpoint selection function $\text{anti-fix}(k) = \{x \mid x \neq k(x)\}$. If we do this, then the resulting scalar open game \mathcal{G} has the same best response relation as a simultaneous move game with 2 utility maximising players and outcome function


$$q(x, y) = \begin{cases} (0, 0) & \text{if } x = y \\ (1, 1) & \text{if } x \neq y \end{cases}$$

This is an *anti-coordination game*. If $X = \{A, B\}$ then the pure Nash equilibria and fixpoints of $\mathbf{B}_{\mathcal{G}}(*, *)$ are (A, B) and (B, A) .

Finally, we interpret $\mathcal{D}_1 = \mathcal{D}_{\text{fix}}$ and $\mathcal{D}_2 = \mathcal{D}_{\text{anti-fix}}$. This is a game in which the first player would like to coordinate with the second, and the second would like to differentiate from the first. This has the same best response relation as *matching pennies*, the game with outcome function

$$q(x, y) = \begin{cases} (1, 0) & \text{if } x = y \\ (0, 1) & \text{if } x \neq y \end{cases}$$

This game has no Nash equilibria in pure strategies.

Figure 3. Definition of $\mathcal{D}_{X,Y}^{\Delta}$

8 Sequential games

In the previous section we showed that a simultaneous move game can be represented as an open game using monoidal products of decisions. In this section we will represent sequential games, in which players can observe previous actions of other players before making their choice.

We focus on the sub-class of finite *sequential games* from [4] introduced in Section 2, which are the finite extensive-form games of perfect information in which at each stage a different player chooses, and the player choosing and the set of possible choices are determined only by the stage number and not the previous moves. That is to say, distinct players $1, \dots, n$ sequentially make choices from sets X_1, \dots, X_n , with each player observing every previous move. Relaxing each of these restrictions is possible but requires more work (generally, defining additional composition operators on open games, such as those in [8]), and so we focus on this sub-class for simplicity.

Recall from Definition 4.1 that the choice of an element of Y after observing an element of X by a utility-maximising agent is modelled by the open game $\mathcal{D}_{X,Y} : (X, 1) \rightarrow (Y, \mathbb{R})$ defined by

- $\Sigma_{\mathcal{D}_{X,Y}} = X \rightarrow Y$
- $\mathbf{P}_{\mathcal{D}_{X,Y}}(\sigma, x) = \sigma(x)$
- $\mathbf{C}_{\mathcal{D}_{X,Y}}(\sigma, x, r) = *$
- $(\sigma, \sigma') \in \mathbf{B}_{\mathcal{D}_{X,Y}}(x, k)$ iff $\sigma'(x) \in \arg \max k$

The basic element of a sequential game is the open game $\mathcal{D}_{X,Y}^{\Delta} : (X, 1) \rightarrow (X \times Y, \mathbb{R})$ denoted by the string diagram in figure 3. Algebraically, this is $(\text{id}_{(X,1)} \otimes \mathcal{D}_{X,Y}) \circ (\Delta_X, 1)$, where $(\Delta_X, 1) : (X, 1) \rightarrow (X \times X, 1)$ is the lifting of the copying function $X \rightarrow X \times X$. By applying the definitions of the composition operators \circ and \otimes , the reader can verify that $\mathcal{D}_{X,Y}^{\Delta}$ is concretely given as follows, up to natural isomorphism:

- $\Sigma_{\mathcal{D}_{X,Y}^{\Delta}} = X \rightarrow Y$
- $\mathbf{P}_{\mathcal{D}_{X,Y}^{\Delta}}(\sigma, x) = (x, \sigma(x))$
- $\mathbf{C}_{\mathcal{D}_{X,Y}^{\Delta}}(\sigma, x, r) = *$
- $(\sigma, \sigma') \in \mathbf{B}_{\mathcal{D}_{X,Y}^{\Delta}}(x, k)$, where $k : X \times Y \rightarrow \mathbb{R}$, iff $\sigma'(x) \in \arg \max_{y:Y} k(x, y)$

Definition 8.1. Let X_1, \dots, X_n be a sequence of sets. We recursively define a sequence of open games

$$\mathcal{G}_i : (1, 1) \rightarrow \left(\prod_{j=1}^i X_j, \mathbb{R}^i \right)$$

as follows. The base case is $\mathcal{G}_0 = \text{id}_{(1,1)} : (1, 1) \rightarrow (1, 1)$. In the recursive step, \mathcal{G}_{i+1} is defined in terms of \mathcal{G}_i and $\mathcal{D}_{\prod_{j=1}^i X_j, X_{i+1}}$ by the string diagram in Figure 4.




Figure 4. Recursive step of definition 8.1




Figure 5. Example sequential open game with 2 players

Theorem 8.2. Let X_1, \dots, X_n be a sequence of sets and $q : \prod_{i=1}^n X_i \rightarrow \mathbb{R}^n$. Let \mathcal{G}_n be defined as in Definition 8.1. Then

$$\Sigma_{\mathcal{G}_n} = \prod_{i=1}^n \left(\prod_{j=1}^{i-1} X_j \rightarrow X_i \right)$$

is the set of strategy profiles of the n -player sequential game with outcome function q , and $\mathbf{B}_{\mathcal{G}_n}(*, q)$ is its best response relation.

For example, a 2-player sequential game with outcome function $q : X \times Y \rightarrow \mathbb{R}^2$ corresponds to the open game depicted in Figure 5. This game has a set of strategy profiles $\Sigma = X \times (X \rightarrow Y)$, and the best response relation $\mathbf{B} \subseteq \Sigma \times \Sigma$ is defined by $(\sigma, \sigma') \in \mathbf{B}$ iff $\sigma'_1 \in \arg \max_{x:X} (q(x, \sigma_2(x)))_1$ and $\sigma'_2(\sigma_1) \in \arg \max_{y:Y} (q(\sigma_1, y))_2$.

We close this section by combining simultaneous and sequential elements. Consider the open game in Figure 6. The game depicts a situation where a player first makes a decision ($\mathcal{D}_{1,X}$). This decision is observed by two players who move simultaneously (\mathcal{D}_{X,Y_1} and \mathcal{D}_{X,Y_2}). As they move simultaneously they cannot observe each others' moves; they do observe the first player's move though. A possible economic story is: two companies use the same input produced by a monopolist, the first player, who sets a price for the input. For example the latter players could be rival car manufacturers, and the first player a monopolist who produces tyres. Upon observing the price, both competitors decide how much to produce. Profits result accordingly.

The example illustrates how compositionality can be applied to economic strategic situations. Modelling economic interactions compositionally is very natural because the object itself is a composition of elements: A market is composed of competitors, buyers, upstream input providers etc. A situation such as this hybrid simultaneous/sequential game is typically modelled as an extensive form game of imperfect information. We suggest that the representation




Figure 6. Hybrid sequential-simultaneous move game

as an open game, as well as being modular, is a more appropriate representation of the underlying economic situation.

9 Conclusions and future work

The broader goal of the compositional theory of games started in this paper is to bring the full force of compositionality to economic modelling. In this paper, we have made an important initial step by showing that a compositional theory is possible. We have focused only on simultaneous and sequential games, on Nash equilibria as the solution concept, and have considered only pure strategies. Of course, there are many more interesting questions to pursue.

First, an important class of games are repeated games, where players engage in an interaction more than once, and especially infinitely repeated games. See [8] for initial results including the construction of final coalgebras of certain functors on open games which can be used to model infinitely repeated games and thereby bring the powerful concept of bisimulation to bear on infinite games.

Secondly, the solution concept of Nash is built into the definition of an open game. It is important to consider alternatives. One categorically attractive point of view is pursued in [11] where open games are taken as objects of a category and morphisms between them are studied. This is particularly important if we want to define open games, or operators on open games, by universal properties.

Thirdly, applications of game theory very commonly use mixed (probabilistic) strategies, since there are games, such as matching pennies, without a Nash equilibrium in pure strategies. Thus, it is important to consider open games with mixed strategies. This is possible but surprisingly difficult, and requires some heavier category-theoretic tools, and is work in progress. A related extension concerns games of incomplete information where some players do not have access to all relevant informations, for instance other players' utilities. Coalgebras are a natural way to study this extension [15].

Fourthly, an important practical question is how open games can be implemented and solved. Various hardness results in algorithmic game theory [17] state that approximating solutions of arbitrary games is computationally hard but this is usually given for classical and 'monolithic' games such as normal-form games, without exploiting a formally defined composed structure. On the other hand,

solution algorithms for economic models usually combine various numerical and statistical methods like function approximation, root finding, integration and Monte Carlo methods, see for example [23]. One area where compositional game theory might have a significant impact is in systematically (functorially) combining numerical methods in order to exploit the compositional structure of the game to solve it.

Finally, comparison with the recent paper [7] suggests a deep and unexpected connection between game theory and deep learning. The category of open learners from that paper embeds into the category of open games, a fact that has several potential applications such as to strategy learning and generative adversarial networks.

References

- [1] Samson Abramsky. 2005. Abstract scalars, loops, and free traced and strongly compact closed categories. In *Proceedings of CALCO'05 (Lectures notes in computer science)*, Vol. 3629. 1–29.
- [2] Achim Blumensath and Viktor Winschel. 2013. A Compositional Coalgebraic Semantics of Strategic Games. (2013). arXiv:1712.08381.
- [3] Bob Coecke and Aleks Kissinger. 2017. *Picturing quantum processes*. Cambridge University Press.
- [4] Martin Escardó and Paulo Oliva. 2011. Sequential games and optimal strategies. *Proceedings of the Royal Society A* 467 (2011), 1519–1545.
- [5] Martin Escardó and Paulo Oliva. 2012. Computing Nash equilibria of unbounded games. *Proceedings of the Turing centenary conference* (2012).
- [6] Brendan Fong. 2016. *The algebra of open and interconnected systems*. Ph.D. Dissertation, University of Oxford.
- [7] Brendan Fong, David Spivak, and Rémy Tuyéras. 2017. Backprop as functor: A compositional perspective on supervised learning. (2017). arXiv:1711.10455.
- [8] Neil Ghani, Clemens Kupke, Alasdair Lambert, and Fredrik Nordvall Forsberg. 2018. A compositional treatment of iterated open games. (2018). To appear in *Sannellabration!*, special issue of *Theoretical computer science*.
- [9] Jules Hedges. 2016. *Towards compositional game theory*. Ph.D. Dissertation, Queen Mary University of London.
- [10] Jules Hedges. 2017. Coherence for lenses and open games. (2017). arXiv:1704.02230.
- [11] Jules Hedges. 2017. Morphisms of open games. (2017). arXiv:1711.07059.
- [12] Jules Hedges, Paulo Oliva, Evguenia Shprits, Viktor Winschel, and Philipp Zahn. 2017. Higher-order decision theory. In *Algorithmic Decision Theory (Lecture Notes in Artificial Intelligence)*, Jörg Rothe (Ed.), Vol. 10576. Springer, 241–254. https://doi.org/10.1007/978-3-319-67504-6_17
- [13] Jules Hedges, Paulo Oliva, Evguenia Shprits, Viktor Winschel, and Philipp Zahn. 2017. Selection equilibria of higher-order games. In *Practical aspects of declarative languages (Lecture Notes in Computer Science)*, Yuliya Lierler and Walid Taha (Eds.), Vol. 10137. Springer, 136–151.
- [14] Kevin Leyton-Brown and Yoav Shoham. 2008. *Essentials of game theory: a concise, multidisciplinary introduction*. Morgan and Claypool.
- [15] Lawrence S. Moss and Ignacio D. Viglizzo. 2004. Harsanyi Type Spaces and Final Coalgebras Constructed from Satisfied Theories. *Electronic Notes in Theoretical Computer Science* 106 (2004), 279 – 295. Proceedings of the Workshop on Coalgebraic Methods in Computer Science (CMCS).
- [16] Roger B. Myerson. 1991. *Game theory: analysis of conflict*. Cambridge: Harvard University Press USA.
- [17] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. 2007. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA.
- [18] Dusko Pavlovic. 2009. A semantical approach to equilibria and rationality. In *Algebra and coalgebra in computer science (Lectures notes in computer science)*, Vol. 5728. Springer, 317–334.
- [19] Matthew Pickering, Jeremy Gibbons, and Nicolas Wu. 2017. Profunctor optics: Modular data accessors. *The art, science and engineering of programming* 1, 2 (2017).
- [20] Peter Selinger. 2011. A survey of graphical languages for monoidal categories. In *New structures for physics*, Bob Coecke (Ed.), Springer, 289–355.
- [21] Michael Shulman. 2010. Constructing symmetric monoidal bicategories. (2010). arXiv:1004.0993.
- [22] Jan Willems. 2007. The behavioural approach to open and interconnected systems. *IEEE Control Systems* 27, 6 (2007), 46–99.
- [23] Viktor Winschel and Markus Krätzig. 2010. Solving, Estimating, and Selecting Nonlinear Dynamic Models Without the Curse of Dimensionality. *Econometrica* 78, 2 (2010), 803–821.