

Choice Structures in Games

Paolo Galeazzi and Johannes Marti

April 23, 2023

Abstract

Following the decision-theoretic approach to game theory, we extend the analysis of [EW96] and [DT08] from hierarchies of preference relations to hierarchies of choice functions. We then construct the universal choice structure containing all these choice hierarchies, and show how the universal preference structure from [DT08] is embedded in it.

1 Introduction

The present work focuses on the foundations of that part of the theory of games that deals with the players' interactive beliefs and subjective preferences. We work here in the broadest setup, where the players' beliefs and preferences are expressed in terms of choice functions with no further properties assumed. This generalizes the existing approaches in the literature, where the players' beliefs and preferences are represented by probability and utility functions or by preference relations. We show that even in our general setup it is possible to build a universal type structure - which we call universal choice structure - containing all the coherent hierarchies of choice functions. By being universal, we mean that such structure is terminal, in the sense of category theory, and complete, in the sense of [Bra03]. Furthermore, the universal choice structure is non-redundant (in the sense of [MZ85]) and topology-free (in the sense of [HS98]). All these notions are discussed in more detail in Section 5.

The motivation for the present work comes from the limitations imposed by employing order relations to represent the players' preferences, as already recognized by [Eps97]. Generalizations of probabilistic type structures, called preference structures, have been introduced in [EW96, DT08] and [GHL16], where types are defined as hierarchies of interactive preference relations rather than as hierarchies of interactive probabilistic beliefs. However, not all important decision criteria are representable by order relations and, consequently, preference structures are of no help to the study of the behavioral implications when rationality is defined in terms of such criteria. A famous example is regret minimization, which violates the principle of independence of irrelevant alternatives and hence transitivity too (see Section 1.2 below). As a consequence, regret-minimizing preferences cannot be represented by order relations and have been axiomatized by means of choice functions (see [Hay08] and [Sto11]).

The present paper aims to build the interactive epistemic structures required to take into account cases of decision criteria like regret minimization, where preference orders are insufficient to represent the players' rationality. Such motivation justifies the shift from hierarchies of interactive preference orders to hierarchies of interactive choice functions.

In addition to the conceptual motivation of this work, choice structures turn out to be the most general models for interactive epistemology and epistemic game theory introduced so far. As preference orders are special cases of choice functions, the universal choice structure embeds the universal preference structure, which in turn embeds the universal probabilistic type structure.

On a more technical side, we prove the formal results of this paper using notions from the theory of coalgebras and category theory. This allows for a modular approach, which may also shed new light on the construction of both the universal probabilistic type structure and the universal preference structure.

The paper is structured as follows. The following two Subsections 1.1 and 1.2 survey the related literature and provide a motivational example for this work, respectively. Section 2 introduces the setting and the preliminary notions that are used in Section 3 to construct hierarchies of choice functions and the universal choice structure. Section 4 shows how the universal preference structure is embedded into the universal choice structure, and Section 5 concludes.

1.1 Related literature

As already pointed out above, [EW96] generalize the notion of a type from a hierarchy of probabilistic beliefs to a hierarchy of interactive preference relations and introduce preference structures accordingly. [Che10] then proves that any structure à la [EW96] can be embedded into their preference structure by a morphism which is unique. Starting from different premises than [EW96], [DT08] also constructs a preference structure which he shows to be universal and non-redundant. In the same spirit as [EW96] and [DT08], [GHL16] show the existence of the universal structure for a larger class of preferences than in [EW96]. Our results are thus a further generalization from preference structures to choice structures tout court, in order to be even more liberal about the players' rationality and decision criteria.

Other related work, more for the formal techniques employed than in the spirit, is [Hei14]. There, the author provides the construction of a universal type structure with unawareness by means of tools from category theory. Different from our case, however, for his construction to succeed it suffices to prove that his type structures with unawareness are coalgebras for an appropriately defined functor, for which the results of [Hei14] directly follow from those in [MV04, Vig05]. In our case, instead, a more extensive use of category theory is necessary, as our choice structures are coalgebras for a new different functor, whose categorical properties have not been investigated yet.

This last observation also relates our results to the work by [MV04, Vig05], in that they are the first to formulate classic probabilistic type structures as coalgebras for the probability measure functor and to show that the universal type structure is the terminal coalgebra in the category of probabilistic type structures appropriately defined. We show here that the terminal coalgebra also exists for choice structures.

1.2 A motivational example

Consider the following game, that Ida is playing as the row player and Joe as the column player.

	<i>l</i>	<i>r</i>
<i>u</i>	5;1	0;0
<i>m</i>	3;2	0;1
<i>c</i>	1;1	3;0
<i>d</i>	1;2	2;3

Ida has two dominated actions: *m* and *d*. Specifically, the mixed action $pu + (1 - p)c$ strongly dominates *m* for $p \in (1/2, 1)$ and *d* for $p \in (0, 1/3)$. Therefore, no probabilistic belief on Joe's actions can justify the choice of either *m* or *d* in terms of expected utility maximization: For every probabilistic belief of Ida, either *u* or *c* will have higher expected utility than *m* as well as *d*. This is evident from Figure 1a, where the horizontal axis represents the probability of Joe playing *r* and the vertical axis represents Ida's expected utility. Actions that are consistent with the players' rationality and common belief in rationality are called rationalizable. In the example, the only action profile that is rationalizable with expected utility maximization is (u, l) , since action *r* is no longer rational when *d* is eliminated.

When, in the wake of decision-theoretic developments, classic probabilistic beliefs are generalized to possibly non-probabilistic beliefs (e.g. [GS89, Sch89]), the game-theoretic notions of rationality, dominance and equilibrium have to be reconsidered accordingly. Examples of these advancements in game theory can be found e.g. in [Kli96, Lo96, Mar00, KU05, RS10, HP12, BCVMM15, Tro19]. For the sake of explanation, let us take the case where beliefs are represented



Figure 1: On the left, the representation of Ida’s expected utilities associated with the actions in the game above. On the right, the representation of Ida’s expected negative regrets associated with the same actions.

by sets of probability distributions as in the multiple-prior (MP, henceforth) model of [GS89]. In this case, each action is associated with a set of expected utilities, and a natural notion of rational choice consists in picking an action with the highest minimal expected utility. When maximin expected utility is substituted for expected utility maximization as the notion of rationality in the presence of non-probabilistic beliefs, a question naturally arises: How is dominance defined in this setting?

An answer is offered by [Eps97], who establishes a correspondence between iterated elimination of MP-dominated actions and MP-rationalizability. To exemplify, consider again the game above between Ida and Joe, and suppose that both players are expected utility maximizers: What are then the behavioral implications of rationality and common belief in rationality? Figure 1a shows that actions u, c and d are justifiable for Ida: u and c are still best replies to some probabilistic belief, while d is now a possible best reply to some non-probabilistic belief, e.g., the set of probability distributions that coincides with the simplex over Joe’s actions. The MP-rationalizable action profiles are therefore the members of the Cartesian product $\{u, c, d\} \times \{l, r\}$.

The results about MP-dominance and MP-rationalizability in [Eps97] are based on preference structures, i.e., type structures whose elements consist in hierarchies of interactive, reflexive and transitive preference relations over Savage-style acts. The space of all coherent preference hierarchies, i.e., the universal preference structure, is thus foundational to the results about iterated dominance and rationalizability for decision criteria that are representable by reflexive and transitive preference relations over acts, such as maximin expected utility and all other noteworthy criteria in [Eps97]. As mentioned above, universal preference structures have been constructed by both [EW96] and [DT08].

Once multiple decision criteria are introduced for single-agent problems, it is also natural to think of situations where different individuals adhere to different criteria in playing games. In such cases, we may have for instance Ida playing the game and choosing actions according to criterion A, while Joe is making his choices according to criterion B. Uncertainty about the opponent’s criterion therefore enters the players’ epistemic state and spreads to higher-order levels too: Ida may be uncertain about Joe’s criterion and about Joe’s uncertainty relative to her criterion, and so on. Preference structures provide a formal framework to express such interactive higher-order uncertainty about the players’ decision criteria.

Consider the game above again, but now suppose that both Ida and Joe are regret minimizers. Figure 1b helps picture the situation, where each action is plotted in terms of its expected negative regret. Taking advantage of the fact that the minimization of the maximal (positive) regret is equivalent to maximin negative regret, Figure 1b shows that action m is now a possible best reply (e.g., to the set of probability distributions coinciding with the simplex over Joe’s actions), whereas



Figure 2: On the left, the representation of Ida’s expected negative regrets associated with the actions in the game above. On the right, the representation of Ida’s expected negative regrets associated with the actions in the game above, when action u is no longer available.

action d is no longer a best reply to any possible belief of Ida. Joe would consequently not play action r , and the only regret-rationalizable profile is thus (u, l) .

As already pointed out, however, regret minimization violates the independence of irrelevant alternatives and is therefore not representable by a preference relation. To see it, we can make use of the game above again. Suppose for instance that Ida is a regret minimizer and her belief is represented by a convex compact set of probability distributions assigning action r a lower probability of 0.25 and an upper probability of 1 (see Figure 2). For this specific belief, Ida is then indifferent between actions u, m and c , when she can choose among her four original actions (Figure 2a). However, when action u is no longer available, Ida will go for action c (Figure 2b): actions m and c are no longer equivalent, in violation of the independence of irrelevant alternatives. This choice pattern cannot be encoded by a preference order, and a choice function C such that

$$C(\{u, m, c, d\}) = \{u, m, c\} \quad C(\{m, c, d\}) = \{c\}$$

has instead to be employed. These cases never arise for maxmin expected utility, in that the relative order among possible alternatives does not change when actions are added or removed. This is essentially the reason why criteria such as maxmin expected utility can be axiomatized by preference orders ([GS89]), while context-dependent criteria like regret minimization require the use of choice functions ([Hay08] and [Sto11]). In interactive contexts, this motivates the generalization from preference structures to choice structures.

2 Preliminaries

In this section we introduce the mathematical notions that will be employed in game-theoretic contexts in the following sections.

2.1 Uncertainty spaces

We are working with *uncertainty spaces*, or simply *spaces*, (X, \mathcal{B}) , where X is any set, whose elements are called *states*, and $\mathcal{B} \subseteq \mathbb{P}X$ is an algebra of subsets of X , where we denote by $\mathbb{P}X$ the powerset of X . The elements of \mathcal{B} are called *measurable sets* or *events*.¹ We often write just X for an uncertainty space (X, \mathcal{B}) . Any set X can be considered as a trivial uncertainty space (X, \mathcal{B}) in which every subset is measurable, meaning that $\mathcal{B} = \mathbb{P}X$ is the *discrete algebra* containing all subsets of X .

¹That \mathcal{B} is an algebra means that it is closed under finite intersections and complements. Thus, an uncertainty space (X, \mathcal{B}) is almost a measurable space, but without requiring closure under countable intersections.

A morphism φ from an uncertainty space X to an uncertainty space Y is any function $\varphi : X \rightarrow Y$ that is *measurable*, meaning that $\varphi^{-1}[E] = \{x \in X \mid \varphi(x) \in E\}$ is measurable in X whenever E is measurable in Y . For every uncertainty space $X = (X, \mathcal{B})$ we use id_X to denote the measurable identity function on X , that is, $\text{id}_X : X \rightarrow X, x \mapsto x$.

Two uncertainty spaces X and Y are *isomorphic*, written as $X \simeq Y$, if there is a bijective function $\varphi : X \rightarrow Y$ such that both φ and its inverse $\varphi^{-1} : Y \rightarrow X$ are measurable, and $\varphi \circ \varphi^{-1} = \text{id}_Y$ and $\varphi^{-1} \circ \varphi = \text{id}_X$.

Given two uncertainty spaces X and Y we can define their product to be the uncertainty space $X \times Y$ whose states are all pairs (x, y) where the first component is a state in X and the second component is a state in Y . The measurable sets of states are generated by taking finite unions and complements of cylinders, that is, sets of the form $U \times Y$ and $X \times V$ for measurable U in X and V measurable in Y . It is clear that with this algebra the projections $\pi_1 : X \times Y \rightarrow X, (x, y) \mapsto x$ and $\pi_2 : X \times Y \rightarrow Y, (x, y) \mapsto y$ are measurable functions. Moreover given two measurable functions $\varphi : X \rightarrow X'$ and $\psi : Y \rightarrow Y'$ we use $\varphi \times \psi : X \times Y \rightarrow X' \times Y'$ to denote the measurable function, which is defined such that $(\varphi \times \psi)(x, y) = (\varphi(x), \psi(y))$ for all $(x, y) \in X \times Y$. It is not hard to check that this $\varphi \times \psi$ is indeed measurable.

2.2 Savage acts

Fix a non-empty set Z of *outcomes*. A *Savage act*, or just *act*, for an uncertainty space X is a measurable finite step function $f : X \rightarrow Z$. That f is a finite step function means that its range $f[X] = \{f(x) \in Z \mid x \in X\}$ is finite. We also assume that the set Z carries the discrete algebra in which all sets are measurable. Using that the measurable sets in X are closed under finite unions, one easily checks that a finite step function $f : X \rightarrow Z$ is measurable precisely if $f^{-1}[\{z\}]$ is measurable for all $z \in Z$. We write FX for the set of all acts for some uncertainty space X . For every measurable function $\varphi : X \rightarrow Y$ we obtain a function $F\varphi : FY \rightarrow FX, f \mapsto f \circ \varphi$.

Our notion of a Savage act, where Z is any set and $f : X \rightarrow Z$ is a finite step function, is then compatible with the settings from both [DT08] and [Sto11]. In the work on preference structures from [DT08], it is assumed that Z is finite and thus all acts are automatically finite step functions. In the framework from [Sto11], acts are measurable finite step functions $f : X \rightarrow \Delta Z$, where ΔZ is the set of all probability distributions with finite support over a set Z of outcomes. This approach can be recovered in our setting by instantiating the set Z with ΔZ .

In the following we are making the assumption that Z is finite, following [DT08]. However, our results on the existence of the universal choice structure hold also for arbitrary, possibly infinite, sets Z . We address this matter in more detail with Remark 3 in Section D.2 of the appendix.

2.3 Choice functions

A *choice function* over a set X is a function C that maps every finite subset $F \subseteq X$ to one of its subsets $C(F) \subseteq F$. For any set X we write $\mathcal{C}X$ for the set of all choice functions over X . Even if X is just a set, without a notion of measurable subset, we take $\mathcal{C}X$ to be the uncertainty space in which a set is measurable if it can be generated by taking finite intersections and complements of sets of the form

$$B_L^K = \{C \in \mathcal{C}X \mid C(K) \subseteq L\},$$

for some finite $K, L \subseteq X$ with $L \subseteq K$.

Given any function $f : X \rightarrow Y$ we define the function $\mathcal{C}f : \mathcal{C}Y \rightarrow \mathcal{C}X$ by setting $\mathcal{C}f(C) = C^f$, where C^f is the choice function mapping a finite $K \subseteq X$ to

$$C^f(K) = f^{-1}[C(f[K])] \cap K.$$

We use $f[K] \subseteq Y$ to denote the direct image $f[K] = \{f(x) \in Y \mid x \in K\}$ of K . One can show that the function $\mathcal{C}f : \mathcal{C}Y \rightarrow \mathcal{C}X$ is measurable. To see this one first checks that for all measurable

sets of the form B_L^K with $K \subseteq L$

$$\begin{aligned} (\mathcal{C}f)^{-1}[B_L^K] &= \{C \in \mathcal{C}Y \mid f^{-1}[C(f[K])] \cap K \subseteq L\} \\ &= \{C \in \mathcal{C}Y \mid C(f[K]) \subseteq \hat{L}\} = B_{\hat{L}}^{f[K]}, \end{aligned}$$

where $\hat{L} = \{y \in f[K] \mid f^{-1}[\{y\}] \cap K \subseteq L\}$. This then extends to arbitrary measurable sets because intersections and complements are preserved under taking inverse images.

2.4 Choice functions over Savage acts

We then consider choice functions over Savage acts for some uncertainty space X . For every uncertainty space X we define the uncertainty space $\Gamma X = \mathcal{C}FX$ to be the uncertainty space of all choice functions over Savage acts for X . Moreover, for every measurable function $\varphi : X \rightarrow Y$ we can define the measurable function $\Gamma\varphi = \mathcal{C}F\varphi : \Gamma X \rightarrow \Gamma Y$. By unfolding the definitions of \mathcal{C} and F , we can describe the choice function $\Gamma\varphi(C) = C^\varphi$ for any $C \in \Gamma X$ more concretely: It maps a finite set of Savage acts $K \subseteq FY$ to the set

$$C^\varphi(K) = \{f \in K \mid f \circ \varphi \in C(\{g \circ \varphi \mid g \in K\})\}.$$

3 Choice structures

In this section, the notions introduced above are applied to the game-theoretic context that we are interested in. To keep things notationally simple, we focus here on interactive situations with only two players, Ida and Joe. It is straightforward to adapt our setting to more than two players, but this would introduce additional notational complications.

3.1 Choice structures

As we are working with two-player games, the basic uncertainty of Ida is just over the fixed finite set A_j of Joe's actions and, similarly, the basic uncertainty of Joe is over the fixed finite set of Ida's actions A_i . We thus obtain the following definition of a choice structure:

Definition 1. A *choice structure* is a tuple $\mathcal{X} = (T_i, T_j, \theta_i, \theta_j)$ consisting of:

- uncertainty spaces T_i and T_j of *types* for Ida and Joe, and
- measurable functions $\theta_i : T_i \rightarrow \Gamma(A_j \times T_j)$ and $\theta_j : T_j \rightarrow \Gamma(A_i \times T_i)$.

A *morphism* $\alpha : \mathcal{X} \rightarrow \mathcal{X}'$ from a choice structure $\mathcal{X} = (T_i, T_j, \theta_i, \theta_j)$ to a choice structure $\mathcal{X}' = (T'_i, T'_j, \theta'_i, \theta'_j)$ consists of two measurable functions $\alpha_i : T_i \rightarrow T'_i$ and $\alpha_j : T_j \rightarrow T'_j$ such that

$$\theta'_i \circ \alpha_i = \Gamma(\text{id}_{A_j} \times \alpha_j) \circ \theta_i \text{ and } \theta'_j \circ \alpha_j = \Gamma(\text{id}_{A_i} \times \alpha_i) \circ \theta_j.$$

A *state* in \mathcal{X} is a tuple $(a_i, a_j, t_i, t_j) \in A_i \times A_j \times T_i \times T_j$. We also say that a *state of Ida* is a pair $(a_i, t_i) \in A_i \times T_i$ and a *state of Joe* is a pair $(a_j, t_j) \in A_j \times T_j$. Notice that the actions specified by Ida's choice function $\theta_i(t_i)$ at state (a_i, a_j, t_i, t_j) may be completely unrelated to the action a_i actually played by Ida at that state (and likewise for Joe). In the example from Section 1, for instance, it is possible that at a given state Ida's type prescribes to choose u , whereas Ida's actual choice is d (see Example 1 below for more details). This is a feature that choice structures have in common with preference structures. Without delving into interpretational issues, here we just think of both choice hierarchies and preference hierarchies as expressing choice attitudes or mental states rather than actual behavior. A state of a player therefore consists of an action that she is actually playing and a type that represents her mental attitude.²

²[BDV21] make a similar point in the context of a different framework, where players have beliefs about their own actions.

A *Savage act of Ida* is then a map $f : A_j \times T_j \rightarrow Z$ from states of Joe to outcomes, and similarly for Joe. We hence model the type of a player by her choices between Savage acts, whose outcome depends on the state of the other player, but not on her own state. This modelling assumes that the players do not have uncertainty about their own type. This is different from the setting of [DT08], who allows players to be uncertain about their own type. We discuss this difference in modelling more extensively in Section 5.3.

Example 1. We provide an example of a choice structure $\mathcal{X} = (T_i, T_j, \theta_i, \theta_j)$ for the game from the introduction. In this example Ida has a single type $T_i = \{t_i\}$, while Joe has two possible types $T_j = \{t_{Mm}, t_{EU}\}$. Here, we interpret Ida's type t_i as a regret minimizer with belief represented as in Figure 2, i.e., a convex compact set of probability distributions assigning action r lower probability of 0.25 and upper probability of 1, while Joe's type t_{Mm} is a maximinimizer with full uncertainty, i.e., not excluding any probability distribution over Ida's actions, and Joe's type t_{EU} is an expected utility maximizer assigning probability 1/2 to Ida playing u and 1/2 to Ida playing d . The states of the choice structure are given by the Cartesian set $A_i \times A_j \times T_i \times T_j$, i.e.,

$$\{u, m, c, d\} \times \{l, r\} \times \{t_i\} \times \{t_{Mm}, t_{EU}\}.$$

Ida's states are then members of the set $\{u, m, c, d\} \times \{t_i\}$ and Joe's states are members of the set $\{l, r\} \times \{t_{Mm}, t_{EU}\}$. The set of outcomes Z is naturally given by the outcomes of the game,

$$Z = \{(5, 1), (3, 2), (1, 1), (1, 2), (0, 0), (0, 1), (3, 0), (2, 3)\}.$$

The map θ_i then associates each of Ida's types with a choice function over Savage acts defined on Joe's states. In the running example, such Savage acts are

$$\begin{array}{lll} f_u(l, t) = (5, 1) & f_u(r, t) = (0, 0) & \text{for both } t \in T_j \\ f_m(l, t) = (3, 2) & f_m(r, t) = (0, 1) & \text{for both } t \in T_j \\ f_c(l, t) = (1, 1) & f_c(r, t) = (3, 0) & \text{for both } t \in T_j \\ f_d(l, t) = (1, 2) & f_d(r, t) = (2, 3) & \text{for both } t \in T_j \end{array}$$

Similarly, Joe's Savage acts are the following:

$$\begin{array}{ll} f_l(u, t_i) = (5, 1) & f_r(u, t_i) = (0, 0) \\ f_l(m, t_i) = (3, 2) & f_r(m, t_i) = (0, 1) \\ f_l(c, t_i) = (1, 1) & f_r(c, t_i) = (3, 0) \\ f_l(d, t_i) = (1, 2) & f_r(d, t_i) = (2, 3) \end{array}$$

When Ida's type t_i is a regret minimizer as described above, the choice function $C_i = \theta_i(t_i)$ associated with t_i maps subsets of the set of Savage acts defined above as follows:

$$\begin{array}{ll} C_i(\{f_u, f_m, f_c, f_d\}) = \{f_u, f_m, f_c\} & C_i(\{f_m, f_d\}) = \{f_d\} \\ C_i(\{f_u, f_m, f_c\}) = \{f_u, f_m\} & C_i(\{f_c, f_d\}) = \{f_c\} \\ C_i(\{f_u, f_m, f_d\}) = \{f_u, f_m\} & C_i(\{f_u, f_m\}) = \{f_u\} \\ C_i(\{f_u, f_c, f_d\}) = \{f_u\} & C_i(\{f_u, f_d\}) = \{f_u\} \\ C_i(\{f_m, f_c, f_d\}) = \{f_c\} & C_i(\{f_u, f_c\}) = \{f_u, f_c\} \\ C_i(\{f_m, f_c\}) = \{f_c\} & \end{array}$$

where we dispense with specifying the choice function in trivial cases such as singletons. The definition of a choice structure would also require C_i to be defined on all other subsets of $F(A_j \times T_j)$. For brevity we give the definition of C_i only on subsets of $\{f_u, f_m, f_c, f_d\}$, which are the Savage acts corresponding to actions in the game. As for Joe, we have that type t_{Mm} is associated with the choice function

$$C_{Mm}(\{f_l, f_r\}) = \{f_l\}$$

and type t_{EU} with the choice function

$$C_{EU}(\{f_l, f_r\}) = \{f_l, f_r\}.$$

Again, we omit the definition of the choice functions on sets of Savage acts which do not correspond to actions in the game.

3.2 Choice hierarchies and the universal choice structure

We now introduce hierarchies of choice functions that represent the higher-order attitudes of the players. To this aim we define uncertainty spaces representing the players' n -th order attitudes by a mutual induction on i and j . In the base case we set $\Omega_{i,1} = \Gamma A_j$ and $\Omega_{j,1} = \Gamma A_i$, and for the inductive step $\Omega_{i,n+1} = \Gamma(A_j \times \Omega_{j,n})$ and $\Omega_{j,n+1} = \Gamma(A_i \times \Omega_{i,n})$. The intuition is that the players' first order attitudes are represented by their choices between Savage acts whose outcome depends just on the actual action played by the opponent. Players' $(n+1)$ -th order attitudes are represented by their choices between acts, whose outcome depends on the actual action played by the opponent and the n -th order attitudes of the opponent.

Note that the players' $(n+1)$ -th order attitudes determine their n -th order attitudes. At the first level this means that the agent's choices in $\Omega_{i,1}$ between Savage acts that depend on the opponent's action are the same as her choices between Savage acts in $\Omega_{i,2}$, when they are taken as additionally depending trivially on the opponent's first-order attitudes. This can be made precise with a measurable *coherence morphism* $\xi_{i,1} = \Gamma\pi_1 : \Omega_{i,2} \rightarrow \Omega_{i,1}$, where $\pi_1 : A_j \times \Omega_{j,1} \rightarrow A_j$ is the projection onto the first component. When $o_2 \in \Omega_{i,2}$ represents Ida's second-order attitudes then $\xi_{i,1}(o_2) \in \Omega_{i,1}$ represents her first-order attitudes.

Analogously, we define a coherence morphism for Joe, by setting $\xi_{j,1} = \Gamma\pi_1 : \Omega_{j,2} \rightarrow \Omega_{j,1}$, where $\pi_1 : A_j \times \Omega_{j,1} \rightarrow A_j$. From now on we will not bother with writing every equation explicitly for Ida and Joe. We just write the version for Ida and then write "and similarly for Joe", thereby meaning that the equation also holds with i and j interchanged.

By induction we can extend the idea of coherence to the higher levels. Choices between acts depending on the opponent's action and n -th order attitudes of the opponent are determined by choices between the same acts taken as depending on the opponent's actions and their $(n+1)$ -th order attitudes. Hence, we define by mutual induction

$$\xi_{i,n+1} = \Gamma(\text{id}_{A_j} \times \xi_{j,n}) : \Gamma(A_j \times \Omega_{j,n+1}) \rightarrow \Gamma(A_j \times \Omega_{j,n})$$

and similarly for Joe $\xi_{j,n+1} = \Gamma(\text{id}_{A_i} \times \xi_{i,n})$.

In the limit one can then consider countable sequences $o = (o_1, o_2, \dots, o_n, \dots)$ with $o_n \in \Omega_{i,n}$ for all $n \in \mathbb{N}$. Moreover, we require these sequences to be coherent in the sense that $\xi_{i,n}(o_{n+1}) = o_n$ for all n . One such sequence completely describes a coherent state of Ida's higher-order attitudes at all levels. Let Ω_i be the infinite set of all such coherent sequences. There are projections $\zeta_{i,n} : \Omega_i \rightarrow \Omega_{i,n}$ for every level $n \in \mathbb{N}$. The set Ω_i becomes an uncertainty space when endowed with the algebra generated from all the subsets of the form $(\zeta_{i,n})^{-1}[O_n]$ for $n \in \mathbb{N}$ and measurable $O_n \subseteq \Omega_{i,n}$. All these notions can also be defined analogously for Joe.

In the appendix we prove the central result about this construction, which is that there exist the following isomorphisms:

Theorem 1. $\Omega_i \simeq \Gamma(A_j \times \Omega_j)$ and $\Omega_j \simeq \Gamma(A_i \times \Omega_i)$.

Using the isomorphisms $\mu_i : \Omega_i \rightarrow \Gamma(A_j \times \Omega_j)$ and $\mu_j : \Omega_j \rightarrow \Gamma(A_i \times \Omega_i)$ from Theorem 1, we then define the universal choice structure:

Definition 2. The *universal choice structure* $\mathcal{U} = (\Omega_i, \Omega_j, \mu_i, \mu_j)$ consists of the uncertainty spaces Ω_i and Ω_j of all coherent sequences of choice attitudes, together with the measurable functions $\mu_i : \Omega_i \rightarrow \Gamma(A_j \times \Omega_j)$ and $\mu_j : \Omega_j \rightarrow \Gamma(A_i \times \Omega_i)$ from Theorem 1.

There is a technical difference between our presentation and the approach that is usually taken in the literature, such as for instance [MZ85, BD93, EW96, DT08]. It is common to define the $(n+1)$ -th level $\Omega_{i,n+1}$ as consisting of all pairs $(x, y) \in \Omega_{i,n} \times \Gamma\Omega_{j,n}$ that are coherent in the sense that the attitudes represented by x are consistent with the attitudes represented by y , in a sense similar to our coherence morphisms $\xi_{i,n}$. In the limit Ω_i one then considers sequences $(o_1, o_2, \dots) \in \Omega_i$ such that $o_n \in \Gamma\Omega_{j,n}$ for all $n \in \mathbb{N}$. Our approach is equivalent to this approach, once coherence of the whole infinite sequences is taken into account.

3.3 Universality of the universal choice structure

Every type $t \in T_i$ in any choice structure $\mathcal{X} = (T_i, T_j, \theta_i, \theta_j)$ generates a coherent sequence of attitudes in Ω_i . To see it, let us define first $v_{i,1} = \Gamma\pi_1 \circ \theta_i : T_i \rightarrow \Omega_{i,1}, t \mapsto \Gamma\pi_1(\theta_i(t))$, where $\pi_1 : A_j \times T_j \rightarrow A_j$ is the projection. Similarly we define $v_{j,1} : T_j \rightarrow \Omega_{j,1}$. We can then continue by mutual induction and set

$$\begin{aligned} v_{i,n+1} &= \Gamma(\text{id}_{A_j} \times v_{j,n}) \circ \theta_i : T_i \rightarrow \Omega_{i,n+1} \\ t &\mapsto \Gamma(\text{id}_{A_j} \times v_{j,n})(\theta_i(t)). \end{aligned}$$

Similarly we define $v_{j,n+1} : T_j \rightarrow \Omega_{j,n+1}$.

One can easily verify that $v_{i,n} = \xi_{i,n} \circ v_{i,n+1}$ for all n . Hence, for each $t \in T_i$ the infinite sequence $(v_{i,1}(t), v_{i,2}(t), \dots)$ is coherent and we obtain a measurable map $v_i : T_i \rightarrow \Omega_i$. Similarly we also obtain a measurable map $v_j : T_j \rightarrow \Omega_j$. In the appendix we show that v_i and v_j together define a unique morphism v into the universal choice structure:

Theorem 2. For every choice structure \mathcal{X} there is a unique morphism of choice structures $v : \mathcal{X} \rightarrow \mathcal{U}$ from \mathcal{X} to the universal choice structure \mathcal{U} .

3.4 Non-redundancy of the universal choice structure

We can also show that our universal choice structure is non-redundant. The notion of non-redundancy goes back to [MZ85], who prove that their universal type space has this property. Our definition of non-redundancy is an adaptation of Definition 3 in [DT08].

For the definition of non-redundancy we consider situations in which we want to endow the sets of types T_i and T_j in a choice structure $\mathcal{X} = (T_i, T_j, \theta_i, \theta_j)$ with an algebra that is possibly distinct from the algebra they have as type spaces in \mathcal{X} . To this aim we need to clarify our notation. Recall that we write an uncertainty space X as (X, \mathcal{B}) , where the symbol X is used for both the underlying set of points and the uncertainty space itself, including the algebra \mathcal{B} . Given an algebra $\mathcal{B}' \subseteq \mathbb{P}X$, which might be distinct from \mathcal{B} , we write (X, \mathcal{B}') for the uncertainty space that has the set X as its underlying set and the algebra \mathcal{B}' as its algebra.

Definition 3. Let $\mathcal{X} = (T_i, T_j, \theta_i, \theta_j)$ be some choice structure. Define the *algebras of observable events* $\mathcal{B}_{\mathcal{X},i}^{ob} \subseteq \mathbb{P}T_i$ and $\mathcal{B}_{\mathcal{X},j}^{ob} \subseteq \mathbb{P}T_j$ to be the smallest algebras such that $\mathcal{B}_{\mathcal{X},i}^{ob}$ contains the sets $B_{\mathcal{X},i}^{ob}(K, L)$ for all finite $K, L \subseteq F(A_j \times (T_j, \mathcal{B}_{\mathcal{X},j}^{ob}))$ and $\mathcal{B}_{\mathcal{X},j}^{ob}$ contains the sets $B_{\mathcal{X},j}^{ob}(K, L)$ for all finite $K, L \subseteq F(A_i \times (T_i, \mathcal{B}_{\mathcal{X},i}^{ob}))$, where we write

$$B_{\mathcal{X},i}^{ob}(K, L) = \{t \in T_i \mid \theta_i(t)(K) \subseteq L\} \quad \text{and} \quad B_{\mathcal{X},j}^{ob}(K, L) = \{t \in T_j \mid \theta_j(t)(K) \subseteq L\}.$$

In Appendix D.4.1 we explain more carefully why these algebras of observable events are well-defined.

Definition 4. An algebra $\mathcal{B} \subseteq \mathbb{P}X$ on the set X is *separating* on X if for any $x, y \in X$ with $x \neq y$ there is a $E \in \mathcal{B}$ such that $x \in E$ and $y \notin E$. The choice structure \mathcal{X} is *non-redundant* if $\mathcal{B}_{\mathcal{X},i}^{ob}$ is separating on T_i and $\mathcal{B}_{\mathcal{X},j}^{ob}$ separating on T_j .

In the appendix we prove that a choice structure \mathcal{X} is non-redundant if and only if the unique morphism v from \mathcal{X} to the universal choice structure \mathcal{U} is injective in both components. Proposition 2.5 in [MZ85], Proposition 2 in [Liu09] and Proposition 2 in [DT08] contain analogous results for probabilistic type spaces and preference structures.

Proposition 1. A choice structure \mathcal{X} is non-redundant if and only if both components v_i and v_j of the unique morphism of choice structures v from \mathcal{X} to the universal choice structure \mathcal{U} are injective.

Because the unique morphism from the universal choice structure \mathcal{U} to itself has the identity function in both components, which is injective, it follows immediately from Proposition 1 that \mathcal{U} is non-redundant.

Corollary 1. The universal choice structure \mathcal{U} is non-redundant.

4 Embedding preference structures

In this section we discuss how our hierarchies of choice functions relate to the hierarchies of preference relations introduced in [DT08].

4.1 Di Tillio's preference structures

We start by reviewing the approach by [DT08] in our notation. The fundamental notion of [DT08] is that of a preference relation over a set X . In the following a *preference relation* \preceq over X is a poset, that is, a reflexive, transitive and anti-symmetric relation, over the set X . We require preference relations to be anti-symmetric. This is different from [DT08], who requires preference relations to be just preorders, that means reflexive and transitive, but not necessarily anti-symmetric relations. We justify this apparent loss of generality in Remark 1 below. In most of our arguments anti-symmetry does not play a role and hence they also work for preorders. Our reason to require anti-symmetry is that in the case of preorders the embedding from preference relations into choice functions need not be injective.

Write $\mathcal{P}X$ for the set of all preference relations over the set X . The set $\mathcal{P}X$ can be turned into an uncertainty space by generating the algebra of measurable events from sets of the form $B_{x_1 \preceq x_2} = \{\preceq \in \mathcal{P}X \mid x_1 \preceq x_2\}$ for some $x_1, x_2 \in X$.

Every function $f : X \rightarrow Y$ gives rise to the measurable function $\mathcal{P}f : \mathcal{P}Y \rightarrow \mathcal{P}X$, where a preference relation \preceq over Y maps to the preference relation $\mathcal{P}f(\preceq) = \preceq^f$ over X that is defined by

$$x_1 \preceq^f x_2 \quad \text{iff} \quad f(x_1) \preceq f(x_2). \quad (1)$$

We can then redo all the constructions from Section 3 using \mathcal{P} instead of \mathcal{C} . Let us sketch how this works. One considers preference relations over acts by considering for every uncertainty space X the space $\Pi X = \mathcal{P}FX$. For every measurable function $\varphi : X \rightarrow Y$ we obtain the measurable function $\Pi\varphi = \mathcal{P}F\varphi : \Pi X \rightarrow \Pi Y$ which maps a preference relation \preceq over FX to the preference relation \preceq^φ over FY that is defined such that $f \preceq^\varphi g$ if and only if $f \circ \varphi \preceq g \circ \varphi$.

Lastly, define a *preference structure* to be a tuple $\mathcal{X} = (T_i, T_j, \theta_i, \theta_j)$ where $\theta_i : T_i \rightarrow \Pi(A_j \times T_j)$ and $\theta_j : T_j \rightarrow \Pi(A_i \times T_i)$ are measurable functions. A *morphism* $\alpha : \mathcal{X} \rightarrow \mathcal{X}'$ from a preference structure $\mathcal{X} = (T_i, T_j, \theta_i, \theta_j)$ to a preference structure $\mathcal{X}' = (T'_i, T'_j, \theta'_i, \theta'_j)$ consists of two measurable functions $\alpha_i : T_i \rightarrow T'_i$ and $\alpha_j : T_j \rightarrow T'_j$ such that

$$\theta'_i \circ \alpha_i = \Pi(\text{id}_{A_j} \times \alpha_j) \circ \theta_i \quad \text{and} \quad \theta'_j \circ \alpha_j = \Pi(\text{id}_{A_i} \times \alpha_i) \circ \theta_j.$$

Note that this definition of a preference structure differs from the one given in [DT08] with respect to the assumption of introspection mentioned in Section 3.1.

The universal preference structure $\mathcal{U}' = (\Omega'_i, \Omega'_j, \mu'_i, \mu'_j)$ can be defined from the limits Ω'_i and Ω'_j of sequences $(\Omega'_{i,n}, \xi'_{i,n})_{n \in \mathbb{N}}$ and $(\Omega'_{j,n}, \xi'_{j,n})_{n \in \mathbb{N}}$ analogous to those that are used to approximate the universal choice structure in the previous section. Hence, set $\Omega'_{i,1} = \Pi A_j$, $\Omega'_{j,1} = \Pi A_i$ and then inductively $\Omega'_{i,n+1} = \Pi(A_j \times \Omega'_{j,n})$ and $\Omega'_{j,n+1} = \Pi(A_i \times \Omega'_{i,n})$. The coherence morphism are such that $\xi'_{i,1} = \Pi\pi_1$, $\xi'_{j,1} = \Pi\pi_2$ and inductively $\xi'_{i,n+1} = \Pi(\text{id}_{A_j} \times \xi'_{j,n})$ and $\xi'_{j,n+1} = \Pi(\text{id}_{A_i} \times \xi'_{i,n})$. The existence of suitable μ'_i and μ'_j and universality properties of \mathcal{U}' then follow from a construction that is analogous to the one given in Appendix D, using Π in place of Γ . The properties of Π that are required for this construction to succeed are stated in Section 3 of [DT08].

4.2 Maximization

We use maximization to map preference orders to choice functions. Given a finite set of alternatives, a player with a given preference relation chooses the most preferred alternatives of the set. Formally, this means that given a preference relation $\preceq \in \mathcal{P}X$ over a set X we map it to the choice function $m_X(\preceq) \in \mathcal{C}X$ that assigns to a finite set $K \subseteq X$ the set $m_X(\preceq)(K)$ of its maximal elements, which is defined as follows:

$$m_X(\preceq)(K) = \{m \in K \mid \text{there is no } k \in K \text{ with } m \prec k\},$$

where $x \prec y$ is defined to hold if and only if $x \preceq y$ and not $y \preceq x$. This definition yields a measurable function $m_X : \mathcal{P}X \rightarrow \mathcal{C}X$ for every set X . To prove that it is measurable it suffices to show that for every basic event $B_L^K = \{C \in \mathcal{C}X \mid C(K) \subseteq L\}$ of $\mathcal{C}X$ its preimage $m_X^{-1}[B_L^K]$ is measurable in $\mathcal{P}X$. To see that this is the case first observe that the set of maximal elements of a set K in a preference relation \preceq is a subset of L if and only if for every element $k \in K \setminus L$ there is some $l \in L$ such that $k \preceq l$ and not $l \preceq k$. Thus we can write

$$m_X^{-1}[B_L^K] = \bigcap_{k \in K \setminus L} \bigcup_{l \in L} (B_{k \preceq l} \setminus B_{l \preceq k}).$$

Since K , L and hence also $K \setminus L$ are finite the right side of the above equality is a finite intersection of finite unions of intersections of basic events with complements of basic events.

Remark 1. With our definition of the choice function $m_X(\preceq)$ from the preference relation \preceq we do not lose any generality by restricting to posets instead of preorders. For every preorder \preceq we can define the poset \preceq' by setting $x \preceq' x'$ iff $x = x'$ or, $x \preceq x'$ and not $x' \preceq x$. One can show that \preceq' is anti-symmetric and that the maximal elements of any finite set $K \subseteq X$ in \preceq' are the same as the maximal elements of K in \preceq , meaning that $m_X(\preceq')(K) = m_X(\preceq)(K)$. It follows from this that any choice function that arises from maximization in some preorder also arises from maximization in some poset. Hence, the restriction to posets does not lose any generality in the class of choice behaviors that we can account for.

There are examples in which the preorders \preceq and \preceq' are distinct. For instance, we might take $X = \{x_1, x_2\}$ to be a two element set and \preceq the total relation, in which the two elements of X are equally preferred, meaning that $x_1 \preceq x_2$ and $x_2 \preceq x_1$. Applying the definition of \preceq' from above shows that \preceq' is then the poset in which x_1 and x_2 are incomparable, meaning that neither $x_1 \preceq' x_2$ nor $x_2 \preceq' x_1$. We have then that $m_X(\preceq)(K) = K = m_X(\preceq')(K)$ for every $K \subseteq X$. Hence, there exist two preorders that account for the same choice behavior. We show in the following proposition that this redundancy disappears if one restricts to posets.

The next proposition shows that any two distinct preference relations give rise via maximization to distinct choice functions. We prove this proposition in the appendix. It relies on our assumption that preference relations are anti-symmetric.

Proposition 2. The measurable function $m_X : \mathcal{P}X \rightarrow \mathcal{C}X$ is injective for all sets X .

We can extend maximization to preference relations and choice functions over acts by defining for every space X the measurable function $\lambda_X = m_{FX} : \Pi X \rightarrow \Gamma X$. A crucial property of λ_X is stated in the following proposition, which we prove in the appendix:

Proposition 3. The mapping λ_X is natural in X . This means that for every measurable function $\varphi : X \rightarrow Y$ we have that $\lambda_Y \circ \Pi \varphi = \Gamma \varphi \circ \lambda_X$.

4.3 Embedding preference structures into choice structures

We can use the mappings λ_X to turn preference structures into choice structures. For every preference structure $\mathcal{X} = (T_i, T_j, \theta_i, \theta_j)$ define the choice structure $\lambda(\mathcal{X}) = (T_i, T_j, \lambda_{A_j \times T_j} \circ \theta_i, \lambda_{A_i \times T_i} \circ \theta_j)$. It is an easy consequence of Proposition 3 that this embedding preserves morphisms in the sense that whenever $\chi : \mathcal{X} \rightarrow \mathcal{X}'$ is a morphism of preference structures then the same pair of measurable functions is also a morphism of choice structures $\chi : \lambda(\mathcal{X}) \rightarrow \lambda(\mathcal{X}')$.

It is also possible to characterize the class of choice structures $\lambda(\mathcal{X}) = (T_i, T_j, \lambda_{A_j \times T_j} \circ \theta_i, \lambda_{A_i \times T_i} \circ \theta_j)$ arising from some preference structure $\mathcal{X} = (T_i, T_j, \theta_i, \theta_j)$. To this end, one can use any of the axiomatizations of the class \mathbb{P} of choice functions that arise from the maximization in a poset. Such an axiomatization is given for instance in Theorem 2.9 of [ABM07]. We then have that a choice structure $\mathcal{X}' = (T'_i, T'_j, \theta'_i, \theta'_j)$ is equal to $\lambda(\mathcal{X})$ for some preference structure \mathcal{X} if and only if for all types $t_i \in T'_i$ and $t_j \in T'_j$ the choice functions $\theta'_i(t_i) \in \Gamma(A_j \times T'_j)$ and $\theta'_j(t_j) \in \Gamma(A_i \times T'_i)$ are in \mathbb{P} .

Because λ_X is injective for every space X , it follows that whenever $\lambda(\mathcal{X}) = \lambda(\mathcal{X}')$ for preference structures \mathcal{X} and \mathcal{X}' then $\mathcal{X} = \mathcal{X}'$. Therefore, any difference between preference structures is preserved in choice structures. One can also define an injective embedding of the universal preference structure into the universal choice structure. Consider the unique morphism of choice structures v from $\lambda(\mathcal{U}')$ to \mathcal{U} that exists by Theorem 2. This morphism consists of two measurable functions $v_i : \Omega'_i \rightarrow \Omega_i$ and $v_j : \Omega'_j \rightarrow \Omega_j$ for which we argue in the appendix that it is injective if A_i and A_j are finite. Hence, we obtain the following theorem:

Theorem 3. The uncertainty spaces Ω'_i and Ω'_j of all preference hierarchies are isomorphic to two subspaces of the spaces of all choice hierarchies Ω_i and Ω_j respectively.

5 Discussion

5.1 Universality of the universal choice structure

The universal choice structure that is defined in this paper has all of the properties that have been associated with the term “universality” in the literature. It is terminal in the categorical sense, meaning that for every choice structure there is a unique morphism to the universal one. It is complete in the sense that the measurable functions μ_i and μ_j of the universal choice structure are onto. The universal choice structure is also non-redundant, as we discussed above. In particular, because we are working in a category of coalgebras, completeness and non-redundancy follow directly from terminality.

5.2 Topological assumptions and coherence

Our construction of the universal choice structure is based on hierarchies of choice functions, where the levels in the hierarchies are linked by what we call coherence morphisms. We use this terminology in analogy to the construction of coherent probabilistic hierarchies by the use of marginalization. However, our result is topology-free, in that we do not need to impose topological assumptions on the spaces. This is different from the classic probabilistic setting, where without additional topological assumptions the universal structure can not be built from coherent hierarchies of beliefs [HS98, HS99, Vig05]. Different from [HS98] and [HS99], the construction of the universal structure directly from coherent choice hierarchies in a topology-free setup succeeds in our case because, like in [DT08], all the sets we deal with are endowed with an algebra, not a σ -algebra, of subsets. We therefore do not incur in the same problem as in [HS99], because the space of all the choice hierarchies is endowed with an algebra, not a σ -algebra, of subsets and the extension of the coherent choice hierarchies to such a space is thus unproblematic.

5.3 Introspection

An important difference between the definitions of type structure in the literature is how they account for a player’s knowledge about her own type. Similar to [EW96], we follow an approach in the spirit of [BD93], where a player has no uncertainty about her own type. This treatment is different from the one in [DT08], who allows for players to be uncertain about their own type. We could use an analogous approach by considering choice structures based on just one uncertainty space X , whose elements should be thought of as the equivalent of the states from Section 3.1, and then define a choice structure with measurable functions $\sigma : X \rightarrow A_i \times A_j$ and $\vartheta : X \rightarrow \Gamma X \times \Gamma X$. The first component of $\vartheta(x) = (\vartheta_i(x), \vartheta_j(x))$ specifies the type of Ida, and the second the type of Joe. As for our proofs, this would mean that we would consider coalgebras for the endofunctor in \mathbf{Unc} that maps an uncertainty space X to the product $(A_i \times A_j) \times (\Gamma X \times \Gamma X)$, and with the obvious behavior on arrows.

To avoid confusion it should be mentioned that the approach of [DT08] is different from the one by [MZ85], even though both settings use analogous mathematical structures. The difference is that [MZ85] impose an introspection condition that requires the players to be certain about

their own type, whereas [DT08] does not. As discussed in [Vig05, sec. 7.4], it is unclear how this condition can be formulated in coalgebraic terms for the structures from [MZ85] and [DT08].

5.4 The coalgebraic approach

For the technical proofs in the appendix we make extensive use of concepts from coalgebra, which in turn uses notions from category theory such as the base category, the coalgebraic functor, terminal objects and natural transformations between functors. The coalgebraic approach provides us with a modular language in which to describe the construction of the universal choice structure. This comes with the following advantages: First, it is easier to adapt the theory to a new representation of uncertainty. Because the general construction of the terminal coalgebra is parametric in the functor, it suffices to define the appropriate functor and check that it preserves limits of suitable diagrams. Second, category theory provides an encompassing language that can be used to compare properties between different kinds of type structures, such as the concepts of completeness, non-redundancy and the like. Third, the coalgebraic approach allows to abstract away from notational complications such as for instance the indexing that is involved in representing games with multiple players. When formulating the terminal-sequence construction, the indexing is conveniently taken care of in the base category and the coalgebraic functor.

5.5 Interpretation of the framework

A point deserving further investigation is the interpretation of both choice hierarchies and preference hierarchies. As a strict behavioral interpretation of such hierarchies seems to conflict with the specification of the players' actual behavior at a state, for now we viewed hierarchies as choice attitudes rather than choice behavior. Deeper analysis on this point is left to more philosophical research in the future, while here we have been primarily concerned with building the necessary technical apparatus introducing hierarchies of choice functions in situations of interactive epistemology.

References

- [ABM07] Fuad Aleskerov, Denis Bouyssou, and Bernard Monjardet. *Utility Maximization, Choice and Preference*. Springer-Verlag, Berlin Heidelberg, 2007.
- [Awo06] S. Awodey. *Category Theory*. Oxford University Press, 2006.
- [BCVMM15] Pierpaolo Battigalli, Simone Cerreia-Vioglio, Fabio Maccheroni, and Massimo Marinacci. Self-confirming equilibrium and model uncertainty. *American Economic Review*, 105:2:646–677, 2015.
- [BD93] A. Brandenburger and E. Dekel. Hierarchies of beliefs and common knowledge. *Journal of Economic Theory*, 59:189–198, 1993.
- [BdRV02] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Cambridge University Press, 2002.
- [BDV21] Pierpaolo Battigalli and Nicodemo De Vito. Beliefs, plans, and perceived intentions in dynamic games. *Journal of Economic Theory*, page 105283, 2021.
- [Bra03] Adam Brandenburger. On the existence of a ‘complete’ possibility structure. *Cognitive processes and economic behavior*, pages 30–34, 2003.
- [Che10] Yi-Chun Chen. Universality of the Epstein–Wang type structure. *Games and Economic Behavior*, 68(1):389–402, 2010.
- [DT08] Alfredo Di Tillio. Subjective expected utility in games. *Theoretical Economics*, 3(3), 2008.

- [Eps97] Larry Epstein. Preference, rationalizability and equilibrium. *Journal of Economic Theory*, 73(1):1–29, 1997.
- [EW96] Larry Epstein and Tan Wang. "Beliefs about Beliefs" without Probabilities. *Econometrica*, 64(6):1343–73, 1996.
- [FHMV03] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning About Knowledge*. MIT Press, 2003.
- [GHL16] Jayant Ganguli, Aviad Heifetz, and Byung Soo Lee. Universal interactive preferences. *Journal of Economic Theory*, 162:237–260, 2016.
- [GS89] Itzhak Gilboa and David Schmeidler. Maxmin Expected Utility with Non-unique Prior. *Journal of Mathematical Economics*, 18:141–153, 1989.
- [Hay08] Takashi Hayashi. Regret aversion and opportunity dependence. *Journal of Economic Theory*, 139(1):242–268, March 2008.
- [Hei14] Sander Heinsalu. Universal type structures with unawareness. *Games and Economic Behavior*, 83:255–266, 2014.
- [HP12] Joseph Y. Halpern and Rafael Pass. Iterated Regret Minimization: a New Solution Concept. *Games and Economic Behavior*, 74:1:194–207, 2012.
- [HS98] A. Heifetz and D. Samet. Knowledge spaces with arbitrarily high rank. *Games and Economic Behavior*, 22:260–273, 1998.
- [HS99] Aviad Heifetz and Dov Samet. Coherent beliefs are not always types. *Journal of Mathematical Economics*, 32(4):475–488, 1999.
- [Jac16] Bart Jacobs. *Introduction to Coalgebra*, volume 59 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2016.
- [Kli96] Peter Klibanoff. Uncertainty, decision, and normal form games, 1996. Manuscript.
- [KU05] Atsushi Kajii and Takashi Ui. Incomplete Information Games with Multiple Priors. *The Japanese Economic Review*, 56:332–351, 2005.
- [Lan98] Saunders Mac Lane. *Categories for the Working Mathematician*. Springer, 1998.
- [Liu09] Qingmin Liu. On redundant types and Bayesian formulation of incomplete information. *Journal of Economic Theory*, 144(5):2115–2145, 2009.
- [Lo96] Kin Chung Lo. Equilibrium in beliefs under uncertainty. *Journal of Economic Theory*, 71(2):443 – 484, 1996.
- [Mar00] Massimo Marinacci. Ambiguous games. *Games and Economic Behavior*, 31(2):191–219, 2000.
- [MV04] Lawrence S. Moss and Ignacio D. Viglizzo. Harsanyi type spaces and final coalgebras constructed from satisfied theories. *Electronic Notes in Theoretical Computer Science*, 106:279–295, 2004.
- [MZ85] J. F. Mertens and S. Zamir. Formulation of Bayesian analysis for games with incomplete information. *International Journal of Game Theory*, 14:1–29, 1985.
- [RS10] Ludovic Renou and Karl H. Schlag. Minimax regret and strategic uncertainty. *Journal of Economic Theory*, 145(1):264 – 286, 2010.
- [Rut00] J.J.M.M. Rutten. Universal coalgebra: A theory of systems. *Theoretical Computer Science*, 249(1):3–80, 2000.

- [Sch89] David Schmeidler. Subjective Probability and Expected Utility Without Additivity. *Econometrica*, 57:571–589, 1989.
- [Sto11] Jörg Stoye. Axioms for minimax regret choice correspondences. *Journal of Economic Theory*, 146:2226–2251, 2011.
- [Tro19] Michael Trost. On the equivalence between iterated application of choice rules and common belief of applying these rules. *Games and Economic Behavior*, 116:1–37, 2019.
- [Vig05] Ignacio Darío Viglizzo. *Coalgebras on Measurable Spaces*. PhD thesis, Indiana University, 2005.
- [Wor99] James Worrell. Terminal sequences for accessible endofunctors. *Electronic Notes in Theoretical Computer Science*, 19:24–38, 1999.

A Overview of the appendix

This appendix contains the technical proofs about the constructions of the universal choice structure from Section 3 and the embedding of the universal preference structure into the universal choice structure from Section 4. We view these constructions as instances of more general constructions from the theory of coalgebras. Thereby, we follow the approach laid out in [MV04, Vig05], who show how classical probabilistic type spaces can be seen as coalgebras. The theory of coalgebras itself makes use of notions from category theory. We review all relevant concepts from category theory in Section B of this appendix. For the interested reader, more in-depth expositions of category theory can be found in [Lan98] and [Awo06]. For a treatment of the theory of coalgebras we refer to [Jac16] or the classic introduction [Rut00].

The cornerstones for the results in this appendix are as follows:

Choice structures as coalgebras: The first insight is that choice structures are coalgebras and morphisms of choice structures are coalgebra morphisms. The definition of a coalgebra is parametric in a fixed endofunctor G on a category \mathcal{D} . The category \mathcal{D} is called the “base category” and the functor G is often called the “type functor”, but we call it the “coalgebraic” functor in order to avoid confusion. We explain the definition of a category and the notion of a functor in Sections B.1 and B.2. Most relevant for us are the category \mathbf{Unc} of uncertainty spaces and measurable functions and the functor Γ on \mathbf{Unc} , which we have already implicitly defined in Section 2.4. However, for technical reasons having to do with the fact that choice structures represent the uncertainty of two players, instead of just one, we have to use \mathbf{Unc}^2 as our base category, which essentially captures situations involving two uncertainty spaces at once, instead of just one as in \mathbf{Unc} . One also has to work with the functor Γ^\heartsuit , which is an adaptation of Γ to \mathbf{Unc}^2 that captures the dependence of a player’s uncertainty on the uncertainty of the other player.

Existence of the terminal coalgebra: The next observation is that the universal choice structure corresponds to what is called the terminal coalgebra in the theory of coalgebras. Our aim is then to use a general coalgebraic construction to show the existence of the terminal coalgebra in the category of choice structures. This construction, called the terminal sequence construction and discussed in [Wor99], can be seen as a coalgebraic counterpart of the construction of the universal type space from the space of coherent hierarchies of beliefs. It can be performed in all categories of coalgebras for a given functor and under certain rather mild conditions on the coalgebraic functor it is guaranteed to yield the terminal coalgebra. There are different versions of these sufficient conditions, but mostly they require that the coalgebraic functor preserves the limits of certain diagrams in the base category. We review the required notion of limits in Section B.5. In the case of the functor Γ^\heartsuit , whose coalgebras are choice structures, we can only show the preservation of a limit for a particular kind of diagram, which we call epic countable cochains. In Section D we review the terminal sequence construction to check that the preservation of limits for epic countable cochains suffices for it to yield the terminal coalgebra.

Preservation of limits of countable epic cochains: To guarantee the existence of the terminal coalgebra we then show in Section C.3 that the functor Γ^\heartsuit preserves limits of epic countable cochains. It is relatively easy to see that this reduces to showing that Γ preserves these limits. The functor Γ can be seen as the composition of two simpler functors, the functor F , which maps an uncertainty space X to the set FX of all Savage acts over X , and the functor \mathcal{C} , which maps a set X to the uncertainty space $\mathcal{C}X$ of all choice functions over X . We then show that both F and \mathcal{C} preserve suitable limits such that their composition Γ preserves limits of countable epic cochains. For F this is a relatively easy observation that is already implicitly part of the proofs in [DT08]. That \mathcal{C} preserves limits of the right kind is Theorem 4, which is the central technical contribution of this paper.

Embedding of preference structures: We also use the coalgebraic approach to prove the results from Section 4 about the embedding of the universal preference structure of [DT08] into the universal choice structure. The setting of preference structures can also be reformulated in the coalgebraic terminology simply by replacing the functor \mathcal{C} mentioned above by a functor \mathcal{P} that captures preference relations. The embedding of preference structures into choice structures then follows from the observation that there is a natural transformation from \mathcal{C} to \mathcal{P} that is injective at all components. We discuss the notion of a natural transformation in Section B.6. In Section E we explain how the existence of such a natural transformation from \mathcal{P} to \mathcal{C} gives rise to an embedding of the universal preference structure into the universal choice structure.

B Preliminaries

We start by reviewing the basics from category theory and the theory of coalgebras.

B.1 Categories

A *category* \mathcal{D} is a collection of objects, possibly a proper class, together with a collection of morphisms, written as $f : X \rightarrow Y$, between any two objects X and Y such that for every object X in \mathcal{D} there is an *identity* morphism $\text{id}_X : X \rightarrow X$ and for all objects X, Y and Z , and morphisms $f : X \rightarrow Y$ from X to Y and $g : Y \rightarrow Z$ from Y to Z their *composition* $g \circ f : X \rightarrow Z$ is a morphism in \mathcal{D} from X to Z . The identities and compositions are required to satisfy the identity law that $f \circ \text{id}_X = f = \text{id}_Y \circ f$ for all morphism $f : X \rightarrow Y$ and the associativity law that $(h \circ g) \circ f = h \circ (g \circ f)$ for all $f : X \rightarrow Y, g : Y \rightarrow Z$ and $h : Z \rightarrow U$. For a morphism $f : X \rightarrow Y$ in \mathcal{D} we call the object X the *domain* of f and the object Y the *codomain* of f .

Examples of categories, which we use in this paper, are the category of sets Set , which has all the sets as its objects and functions between them as morphisms, and the category of uncertainty spaces Unc , which has uncertainty spaces as objects and measurable functions as morphisms. These clearly satisfy the identity and composition laws if one uses the standard identity functions as identity morphisms and the standard compositions of functions for the composition. A different kind of example is the category CS of all choice structures as defined in Section 3 with morphisms of choice structures between them. The results of this paper mostly concern the structure of this category CS .

A more technical example of a category is Unc^2 , which has pairs $X = (X_0, X_1)$ of uncertainty spaces X_0 and X_1 as objects and in which the morphisms from $X = (X_0, X_1)$ to $Y = (Y_0, Y_1)$ are pairs of measurable functions $\varphi = (\varphi_0, \varphi_1)$. One can check that this is a category, when one gives the obvious component-wise definition of identities and composition. We use the category Unc^2 as a convenient abstraction for the technical results of this appendix because it concisely represents situations that concern the uncertainty of two players. In an object $X = (X_0, X_1)$ of Unc^2 , the uncertainty space X_0 represents the uncertainty of Ida, whereas X_1 represents the uncertainty of Joe.

B.2 Functors

A *functor* G from a category \mathcal{D} to a category \mathcal{E} is an assignment of an object GX in \mathcal{E} to every object X of \mathcal{D} and an assignment of a morphism $Gf : GX \rightarrow GY$ in \mathcal{E} to every morphism $f : X \rightarrow Y$ in \mathcal{D} that preserves identities and compositions, meaning that $G\text{id}_X = \text{id}_{GX}$ for all X and $G(g \circ f) = Gg \circ Gf$ for all g and f . If G is a functor from a category \mathcal{D} to the same category \mathcal{D} one also calls G an *endofunctor*.

A simple example of a functor is the powerset functor \mathbb{P} from **Set** to **Set**. It maps a set X to its powerset $\mathbb{P}X$ and a function $f : X \rightarrow Y$ to the direct image map $\mathbb{P}f : \mathbb{P}X \rightarrow \mathbb{P}Y$, which is defined such that $\mathbb{P}f(U) = f[U] = \{f(x) \in Y \mid x \in U\}$ for all $U \subseteq X$. The functor \mathbb{P} is an endofunctor on the category of sets.

Another example of an endofunctor in **Set** is the finite distribution functor \mathbb{D} . It maps a set X to the set $\mathbb{D}X$ of all probability distributions over X that have finite support. A function $f : X \rightarrow Y$ is mapped to the function $\mathbb{D}f : \mathbb{D}X \rightarrow \mathbb{D}Y$ which is such that $\mathbb{D}f(\mu) = \mu^f$ is the probability distribution μ^f on Y such that $\mu^f(V) = \sum_x \mu(x)$ for $x \in f^{-1}[V]$.

A *contravariant functor* from \mathcal{D} to \mathcal{E} is a functor from \mathcal{D} to \mathcal{E}^{op} , where \mathcal{E}^{op} is just like \mathcal{E} , but the direction of all morphisms is turned around. When G is a contravariant functor from \mathcal{D} to \mathcal{E} , one usually just thinks of it as turning morphisms around, in the sense that $Gf : GY \rightarrow GX$ for all $f : X \rightarrow Y$. If one wants to emphasize that a functor G from \mathcal{D} to \mathcal{E} is not contravariant, one calls G *covariant*.

An example of a contravariant functor from this paper is the mapping F sending an uncertainty space X to the set FX of Savage acts for X and a measurable function $\varphi : X \rightarrow Y$ to the function $F\varphi : FY \rightarrow FX$. This is a contravariant functor from the category of uncertainty spaces to the category of sets.

Another example of a contravariant functor is the mapping \mathcal{C} from Section 2.3, which sends a set X to the uncertainty space $\mathcal{C}X$ of choice functions over that set and a function $f : X \rightarrow Y$ to the measurable function $\mathcal{C}f : \mathcal{C}Y \rightarrow \mathcal{C}X$. This defines a contravariant functor from the category of sets to the category of uncertainty spaces. It is straightforward to check that it preserves identities. To see that \mathcal{C} preserves the composition of morphisms, one sees from the definitions that this amounts to checking the following equality:

$$f^{-1}[g^{-1}[C(g[f[K]])]] \cap K = f^{-1}[g^{-1}[C(f[g[K]])] \cap g[K]] \cap K,$$

for all finite sets $K \subseteq Y$, and functions $f : X \rightarrow Y$ and $g : Y \rightarrow W$.

The mapping Γ from Section 2.4 is a covariant endofunctor on the category of uncertainty spaces that results from first applying the contravariant functor F from uncertainty spaces to sets and then the contravariant functor \mathcal{C} to get back to uncertainty spaces. Formally, this means that Γ maps an uncertainty space X to the space $\Gamma X = \mathcal{C}FX$ and a measurable function $\varphi : X \rightarrow Y$ to the measurable function $\Gamma\varphi = \mathcal{C}F\varphi : \Gamma X \rightarrow \Gamma Y$. Because the effects of the two contravariant functors F and \mathcal{C} cancel out, this means that Γ is a covariant functor.

The mappings \mathcal{P} and Π from [DT08], which we survey in Section 4.1, can also be seen as functors analogous to \mathcal{C} and Γ . That is, \mathcal{P} is a contravariant functor from **Set** to **Unc** and Π is the covariant endofunctor on **Unc** that results from first applying the contravariant functor F from **Unc** to **Set** and then the functor \mathcal{P} .

B.3 Coalgebras

We now explain how choice structures can be seen as an instance of the more general notion of a coalgebra for an endofunctor in some category.

To define the notion of a coalgebra we need to fix a category \mathcal{D} and an endofunctor G on \mathcal{D} , that is, a functor G from \mathcal{D} to \mathcal{D} . A *G -coalgebra* (X, ξ) is an object X of \mathcal{D} together with a morphism $\xi : X \rightarrow GX$ in \mathcal{D} . The class of all G -coalgebras can be turned into a category by taking as morphisms from a G -coalgebra (X, ξ) to a G -coalgebra (Y, ν) all morphisms $f : X \rightarrow Y$ in the category \mathcal{D} that satisfy the property that $\nu \circ f = Gf \circ \xi$.

A relatively simple example of a category of coalgebras for the powerset functor \mathbb{P} in the category **Set**. One can see that \mathbb{P} -coalgebras are essentially the same as Kripke frames [FHMV03, BdRV02]. For any Kripke frame (W, R) , with a set W and relation $R \subseteq W^2$, we can define a \mathbb{P} -coalgebra (W, α) where $\alpha : W \rightarrow \mathbb{P}W$ is such that $\alpha(w) = R[\{w\}] = \{v \in W \mid (w, v) \in R\}$ for all $w \in W$. Conversely, every \mathbb{P} -coalgebra (W, α) also defines a Kripke frame (W, R) such where $R = \{(w, v) \mid v \in \alpha(w)\}$. One can also verify that the coalgebra morphisms between \mathbb{P} -coalgebras correspond precisely to the morphisms between Kripke frames.

Choice structures for fixed sets A_i and A_j of actions can be defined as coalgebras for a particular functor Γ^\heartsuit over the category \mathbf{Unc}^2 . The functor Γ^\heartsuit maps an object (X_0, X_1) to the object $(\Gamma(A_j \times X_1), \Gamma(A_i \times X_0))$. The change of indices in the components is intentional, as it corresponds to encoding attitudes about the opponent's attitudes. For morphisms, the functor Γ^\heartsuit sends a pair of measurable functions (φ_0, φ_1) from (X_0, X_1) to (Y_0, Y_1) to the pair of measurable functions $(\Gamma(\text{id}_{A_j} \times \varphi_1), \Gamma(\text{id}_{A_i} \times \varphi_0))$ from $(\Gamma(A_j \times X_1), \Gamma(A_i \times X_0))$ to $(\Gamma(A_j \times Y_1), \Gamma(A_i \times Y_0))$.

We can view every choice structure $\mathcal{X} = (T_i, T_j, \theta_i, \theta_j)$, defined according to Definition 1, as a Γ^\heartsuit -coalgebra $((T_i, T_j), (\theta_i, \theta_j))$ on the object (T_i, T_j) of \mathbf{Unc}^2 with the morphism $(\theta_i, \theta_j) : (T_i, T_j) \rightarrow \Gamma^\heartsuit(T_i, T_j)$. Conversely, every Γ^\heartsuit -coalgebra $((X_0, X_1), (\xi_0, \xi_1))$ contains measurable functions $\xi_0 : X_0 \rightarrow \Gamma(A_j \times X_1)$ and $\xi_1 : X_1 \rightarrow \Gamma(A_i \times X_0)$ and thus gives rise to a choice structure (X_0, X_1, ξ_0, ξ_1) . One can also check that the morphisms of choice structures, as defined in Definition 1, correspond precisely to the coalgebra morphisms between Γ^\heartsuit -coalgebras.

Analogous to choice structures, one can view preference structures as coalgebras for a functor. To this aim one considers the functor Π^\heartsuit from \mathbf{Unc}^2 to \mathbf{Unc}^2 that is defined by replacing all occurrences of Γ in the definition of Γ^\heartsuit with Π . This means that Π^\heartsuit sends an object (X_0, X_1) to $(\Pi(A_j \times X_1), \Pi(A_i \times X_0))$ and a morphism (φ_0, φ_1) to $(\Pi(\text{id}_{A_j} \times \varphi_1), \Pi(\text{id}_{A_i} \times \varphi_0))$.

B.4 Isomorphisms and epic or monic morphisms

A morphism $f : X \rightarrow Y$ in a category \mathcal{D} is an *isomorphism* if there exists a morphism $g : Y \rightarrow X$ such that $g \circ f = \text{id}_X$ and $f \circ g = \text{id}_Y$. One writes $X \simeq Y$ in case there exists an isomorphism $f : X \rightarrow Y$. One can check that in the category of sets the isomorphisms are precisely the bijective functions and in the category \mathbf{Unc} they are the bijective measurable functions whose inverse is also measurable, analogously to the homeomorphism between topological spaces. The isomorphisms in \mathbf{Unc}^2 are those pairs (φ_0, φ_1) of morphisms such that both φ_0 and φ_1 are isomorphisms in \mathbf{Unc} .

A morphism $f : X \rightarrow Y$ is *epic* if for all further morphisms $g, h : Y \rightarrow T$ it holds that if $g \circ f = h \circ f$ then already $g = h$. A morphism $f : X \rightarrow Y$ is *monic* if for all further morphisms $g, h : T \rightarrow X$ it holds that if $f \circ g = f \circ h$ then already $g = h$. One can check that in the category of sets epic morphisms are precisely the surjective functions and monic morphisms precisely the injective functions. Similarly, in the category of uncertainty spaces the epic and monic morphisms are the surjective and injective measurable functions, respectively. Epic and monic morphisms in \mathbf{Unc}^2 are epic and monic in both components. That is, (φ_0, φ_1) is epic if both φ_0 and φ_1 are epic, and analogously for monic.

A functor F from \mathcal{D} to \mathcal{E} *preserves* an epic morphism f of \mathcal{D} if Ff is epic in \mathcal{E} , and F *preserves* a monic morphism f if Ff is monic. Note that, because an epic morphism in \mathcal{E}^{op} is monic in \mathcal{E} , a contravariant functor F from \mathcal{D} to \mathcal{E} preserves an epic morphism f if Ff is monic in \mathcal{E} , and analogously for the preservation of monic morphisms.

We show in Sections C.1 and C.2 that Γ preserves all epic morphisms and that it preserves all monic morphisms $\varphi : X \rightarrow Y$ for which Y is a discrete space.

B.5 Terminal objects, products and limits of chains

We make use of three different instances of the notion of a limit: terminal objects, products and limits of chains. We define these three concepts separately in the following, but to the interested reader it should be noted that they are all instances of the more general concept of a limit for a diagram, which is discussed in all of the texts on category theory cited above.

A *terminal object* in a category \mathcal{D} is any object \top with the property that for every object T of \mathcal{D} there exists a unique morphism $!_T : T \rightarrow \top$. It follows directly from this definition that any two terminal objects in some category need to be isomorphic. To see this assume that we have a second terminal object \top' . Because \top' is terminal we have the morphism $!_{\top'} : \top' \rightarrow \top$. Because \top is universal there is the morphism $!_{\top} : \top' \rightarrow \top$. We have that $\text{id}_{\top} = !_\top = !_{\top'} \circ !_{\top'}$ because $!_{\top}$ is the unique morphism from \top to \top . A similar argument shows that $\text{id}_{\top'} = !_{\top'} \circ !_\top$. As a consequence one can assume that if a category has a terminal object then it is unique, at least if one only cares about the identity of objects up-to isomorphism.

The terminal object in the category Unc can be defined as the uncertainty space $\top = (\{\star\}, \{\emptyset, \{\star\}\})$ which contains just one point. The terminal object in the category Unc^2 can be defined componentwise as the object (\top, \top) where \top is the terminal object of Unc . A more interesting example of a terminal object is the universal choice structure. The universality property from Theorem 2 expresses precisely that the universal choice structure is the terminal object in the category of choice structures.

Given two objects X and Y of a category \mathcal{D} , a *product* of X and Y is an object of \mathcal{D} , written as $X \times Y$, together with two morphisms $\pi_1 : X \times Y \rightarrow X$ and $\pi_2 : X \times Y \rightarrow Y$, which have the following universal property: For every object T and morphisms $f : T \rightarrow X$ and $g : T \rightarrow Y$ there is a morphism $u : T \rightarrow X \times Y$, often written as (f, g) , that is unique for having the property that $f = \pi_1 \circ u$ and $g = \pi_2 \circ u$.

Similar as for terminal objects, one can see that any two objects with this universal property, relative to fixed X and Y , need to be isomorphic. Hence, one often speaks of the product of X and Y as if it was unique.

In the category Unc one can define the product $X \times Y$ of two uncertainty spaces X and Y to be an uncertainty space whose states are all pairs (x, y) , where x is a state of X and y is a state of Y . The measurable sets of states are generated by taking finite unions and complements of cylinders, that is, sets of the form $U \times Y$ and $X \times V$ for measurable U in X and V measurable in Y . It is clear that with this algebra the projections $\pi_1 : X \times Y \rightarrow X, (x, y) \mapsto x$ and $\pi_2 : X \times Y \rightarrow Y, (x, y) \mapsto y$ are measurable functions. Moreover, one can check that this definition of the product has the universal property that for each space T together with measurable functions $\varphi : T \rightarrow X$ and $\psi : T \rightarrow Y$ there is a unique measurable function $\mu : T \rightarrow X \times Y, t \mapsto (\varphi(t), \psi(t))$ such that $\varphi = \pi_1 \circ \mu$ and $\psi = \pi_2 \circ \mu$.

We need a further piece of notation related to the product: Given $f : X \rightarrow Z$ and $g : Y \rightarrow U$ we write $f \times g : X \times Y \rightarrow Z \times U$ for the morphism $f \times g = (f \circ \pi_1, g \circ \pi_2)$. It is easy to check that this definition generalizes our earlier definition of $\varphi \times \psi$ for measurable functions φ and ψ from Section 2.1.

A crucial categorical notion for the construction of the universal choice structure is that of the limit of a countable cochain. A *countable cochain* $(X_n, f_n)_{n \in \omega}$ in a category \mathcal{D} consists of an object X_n of \mathcal{D} for every natural number $n \in \omega$ and a morphism $f_n : X_{n+1} \rightarrow X_n$ for every $n \in \omega$. Here, and in the following we use ω to denote the set of all natural numbers including 0. The f_n are also called the *coherence morphism* of the cochain.

It is clear that we can compose the morphisms f_n to obtain a morphisms $f_n^m = f_n \circ \dots \circ f_{m-1} : X_m \rightarrow X_n$ for all $n, m \in \omega$ with $n \leq m$, where $f_n^n = \text{id}_{X_n} : X_n \rightarrow X_n$ is just the identity on X_n . We call the countable cochain $(X_n, f_n)_{n \in \omega}$ *epic*, if all the morphisms f_n are epic.

A *limit* of the countable cochain $(X_n, f_n)_{n \in \omega}$ is an object X_ω together with projection morphisms $p_n : X_\omega \rightarrow X_n$ such that $p_n = f_n \circ p_{n+1}$ for all $n \in \omega$ which satisfies the following universal property: For every object T together with morphisms $g_n : T \rightarrow X_n$, satisfying $g_n = f_n \circ g_{n+1}$ for all $n \in \omega$, there is a morphism $u : T \rightarrow X_\omega$ that is unique for having the property that $g_n = p_n \circ u$ for all $n \in \omega$. As for the product and the terminal object, one can show with this universal property that any two limits of a given cochain are isomorphic.

For every cochain $(X_n, \xi_n)_{n \in \omega}$ in the category Unc we can define its limit X_ω to be the uncertainty space that has as its states all sequences $x = (x_0, x_1, \dots, x_n, \dots)$, where x_n is a state from X_n for all $n \in \omega$, which are coherent in the sense that $x_n = \xi_n(x_{n+1})$ for all $n \in \omega$. We can then consider the projections $\zeta_n : X_\omega \rightarrow X_n, x \mapsto x_n$ for each $n \in \omega$. The measurable sets of the limit X_ω are defined to be all sets of the form $\zeta_n^{-1}[E]$ for some $n \in \omega$ and measurable set E of

X_n . It is clear that this definition turns the projections $\zeta_n : X_\omega \rightarrow X_n$ into measurable functions. Moreover, using that all ξ_n are measurable, one can show that the measurable sets defined in this way are closed under finite unions and intersections. To check that this definition of the limit has the universal property observe that for any measurable space T with measurable functions $\varphi_n : T \rightarrow X_n$, such that $\varphi_n = \xi_n \circ \varphi_{n+1}$ for each $n \in \omega$, we can define the unique morphism $\mu : T \rightarrow X_\omega$ such that $\mu(t) = (\varphi_0(t), \varphi_1(t), \dots, \varphi_n(t), \dots)$ for all $t \in T$.

One can also define the limit for every cochain $((X_{0,n}, X_{1,n}), (\xi_{0,n}, \xi_{1,n}))_{n \in \omega}$ in \mathbf{Unc}^2 . This can be done componentwise, meaning that for each of the components of the pairs representing objects and morphisms in \mathbf{Unc}^2 we can just use the construction for \mathbf{Unc} , described in the previous paragraph. For instance the limit itself is just the object $(X_{0,\omega}, X_{1,\omega})$ such that $X_{0,\omega}$ and $X_{1,\omega}$ are the limits of $(X_{0,n}, \xi_{0,n})_{n \in \omega}$ and $(X_{1,n}, \xi_{1,n})_{n \in \omega}$ in \mathbf{Unc} . It is not hard to see that this again has the required universal property.

A functor G from \mathcal{D} to \mathcal{E} *preserves* a limit X_ω of a countable cochain $(X_n, f_n)_{n \in \omega}$ if GX_ω is the limit for the countable cochain $(GX_n, Gf_n)_{n \in \omega}$. The main technical result in this appendix is that the functor Γ preserves limits of countable epic cochains. In the proof of this claim in Section C.3 we investigate the preservation of limits of cochains under the contravariant functors F from \mathbf{Unc} to \mathbf{Set} and \mathcal{C} from \mathbf{Set} to \mathbf{Unc} . For this purpose we need to understand limits of cochains in \mathbf{Set}^{op} . Because morphisms in \mathbf{Set}^{op} are just morphisms from \mathbf{Set} with swapped domain and codomain we have that a countable cochain $(X_n, f_n)_{n \in \omega}$ in \mathbf{Set}^{op} is just what is called a countable chain in \mathbf{Set} , meaning that the X_n are all sets and the $f_n : X_n \rightarrow X_{n+1}$ are functions that now go from X_n to X_{n+1} . Moreover the limit X_ω of this cochain in \mathbf{Set}^{op} is what is actually called a colimit of chain in \mathbf{Set} . The colimit X_ω of a chain $(X_n, f_n)_{n \in \omega}$ in \mathbf{Set} can be described concretely as the disjoint union of all the X_n modulo the equivalence relation that identifies an x_n from X_n with an x_m from X_m if there is some $k \geq n, m$ such that $f_k^n(x_n) = f_k^m(x_m)$. Clearly, we can then define inclusions $p_\omega^n : X_n \rightarrow X_\omega$ for all $n \in \omega$. The universal property of this colimit is just the same as the universal property of limit in \mathbf{Set}^{op} with all morphisms turned around: For every set T with functions $h_n : X_n \rightarrow T$ such that $h_n = h_{n+1} \circ f_n$ for each $n \in \omega$, there is a unique function $u : X_\omega \rightarrow T$ such that $h_n = u \circ p_\omega^n$ for all $n \in \omega$.

B.6 Natural transformations

Let F and G be functors that both go from a category \mathcal{D} to a category \mathcal{E} . Then a *natural transformation* η from F to G is an assignment of a morphism $\eta_X : FX \rightarrow GX$ in \mathcal{E} to every object X in \mathcal{D} such that for all morphisms $f : X \rightarrow Y$ in \mathcal{D} it holds that $Gf \circ \eta_X = \eta_Y \circ Ff$, where η_Y is the morphism that η assigns to the object Y .

An example of a natural transformation from the discrete distribution functor to the powerset functor is mapping a finite distribution to its support set. More formally, we have for every set X a function $\text{supp}_X : \mathbb{D}X \rightarrow \mathbb{P}X$ that maps a discrete distribution μ over X to the support set $\text{supp}_X(\mu)$ of μ . The naturality condition that $\mathbb{P}f \circ \text{supp}_X = \text{supp}_Y \circ \mathbb{D}f$ for all functions $f : X \rightarrow Y$ amounts to the fact that for every discrete distribution μ over X there is some $x \in X$ with $f(x) = y$ and $\mu(x) \neq 0$ if and only if $\mu^f(y) \neq 0$, where $\mu^f = \mathbb{D}f(\mu)$.

An example of natural transformation from this paper is the measurable function $\lambda_X : \Pi X \rightarrow \Gamma X$ that is defined for every uncertainty space X . Proposition 3, which is proven in Section E.3, shows that λ is a natural transformation from the functor Π to the functor Γ .

C Properties of Γ

In this part of the appendix we prove crucial properties of the functor Γ which are needed for the coalgebraic construction of the universal choice structure. In case the reader tries to get an overview of the construction of the universal choice structure, they might skip this section and then refer back as needed when reading Sections D and E.

C.1 Γ preserves epic morphisms

We show that Γ preserves epic morphisms, which in \mathbf{Unc} are just surjective measurable functions. Because Γ is the composition of F and \mathcal{C} it suffices to show that F maps surjective measurable functions to injective functions in \mathbf{Set} and that \mathcal{C} maps injective functions to surjective measurable functions. It is easy to check the former using the definition of F and the cancellation property of epic morphism. That \mathcal{C} maps injective functions to surjective measurable functions is shown in the following lemma:

Lemma 1. If $f : X \rightarrow Y$ is injective then $\mathcal{C}f : \mathcal{C}Y \rightarrow \mathcal{C}X$ is surjective.

Proof. It is easy to see that as f is injective then $f^{-1}[f[U]] = U$ holds for all $U \subseteq X$. To check that $\mathcal{C}f$ is surjective consider any $C \in \mathcal{C}X$. We have to find a $C' \in \mathcal{C}Y$ such that $\mathcal{C}f(C') = C$. Define C' such that $C'(K) = f[C(f^{-1}[K])]$ for all finite $K \subseteq Y$. The following equality then follows for all $L \subseteq X$:

$$\begin{aligned} \mathcal{C}f(C')(L) &= f^{-1}[C'(f[L])] \cap L \\ &= f^{-1}[f[C(f^{-1}[f[L]])]] \cap L \\ &= C(L) \cap L = C(L). \end{aligned}$$

□

C.2 Γ preserves monic morphisms to discrete spaces

We show that Γ preserves monic morphisms, if the codomain has the discrete algebra. We again split the problem into two steps. First we show that F maps every surjective measurable function with a discrete codomain to an injective functions and then we show that \mathcal{C} maps injective functions to surjective measurable functions.

Lemma 2. If the measurable function $\varphi : X \rightarrow Y$ is injective and Y has the discrete algebra then the function $F\varphi : FY \rightarrow FX$ is surjective.

Proof. Pick any $f \in FX$. We want to find a $f' \in FY$ such that $f = f' \circ \varphi$. We define $f'(y) = f(x)$, if there is some $x \in X$ such that $\varphi(x) = y$, and $f'(y) = z'$ for an arbitrary $z' \in Z$ otherwise. This is well-defined because φ is injective. The function f' is trivially measurable because Y has the discrete algebra and $f = f' \circ \varphi$ by definition of f' . □

Lemma 3. If $f : X \rightarrow Y$ is surjective then $\mathcal{C}f : \mathcal{C}Y \rightarrow \mathcal{C}X$ is injective.

Proof. It is easy to see that as f is surjective it holds for all $U \subseteq Y$ that $f[f^{-1}[U]] = U$. To check that $\mathcal{C}f$ is injective consider any two $C, C' \in \mathcal{C}Y$ such that $\mathcal{C}f(C) = \mathcal{C}f(C')$. We have to show that then $C(K) = C'(K)$ for all finite $K \subseteq Y$. Using that $f[f^{-1}[U]] = U$, we can compute:

$$\begin{aligned} \mathcal{C}f(C)(f^{-1}[K]) &= f^{-1}[C(f[f^{-1}[K]])] \cap f^{-1}[K] \\ &= f^{-1}[C(K)] \cap f^{-1}[K] = f^{-1}[C(K) \cap K] = f^{-1}[C(K)] \end{aligned}$$

Similarly we obtain also for C' that $\mathcal{C}f(C')(f^{-1}[K]) = f^{-1}[C'(K)]$. From the assumption that $\mathcal{C}f(C) = \mathcal{C}f(C')$ it follows that $f^{-1}[C(K)] = f^{-1}[C'(K)]$. Because f is surjective it follows that $C(K) = C'(K)$. □

C.3 Γ preserves limits of countable epic cochains

In this section we prove the main technical result of this paper:

Theorem 4. Γ preserves limits of countable epic cochains. This means that whenever we have a limit X_ω , with projections $\zeta_n : X_\omega \rightarrow X_n$ for all $n \in \omega$, of a countable cochain $(X_n, \xi_n)_{n \in \omega}$ in which the coherence morphisms ξ_n are epic for all n , then ΓX_ω , with projections $\Gamma \zeta_n : \Gamma X_\omega \rightarrow \Gamma X_n$,

satisfies the universal property of the limit of the cochain $(\Gamma X_n, \Gamma \xi_n)_{n \in \omega}$, i.e., for every uncertainty space T together with measurable functions $\varphi_n : T \rightarrow \Gamma X_n$ for all $n \in \omega$ with $\varphi_n = \Gamma \xi_n \circ \varphi_{n+1}$ there is a unique measurable function $\mu : T \rightarrow \Gamma X_\omega$ such that $\varphi_n = \Gamma \zeta_n \circ \mu$ for all $n \in \omega$.

For the proofs from later sections in this appendix we need the following corollary of Theorem 4.

Corollary 2. Γ^\heartsuit preserves limits of countable epic cochains.

We split the proof of Theorem 4 into two steps. First, we show that F turns limits of epic cochains in Unc into colimits of chains in Set . Second, we show that \mathcal{C} turns colimits of chains in Set into limits of cochains in Unc . It follows that Γ preserves limits of epic cochains because Γ is defined as the composition of F and \mathcal{C} .

To prove the preservation result about F we need two technical lemmas. The first relates acts and measurable functions. It crucially relies on our assumptions that acts are finite step functions.

Lemma 4. Let $\varphi : X \rightarrow Y$ be a measurable function and $f \in FX$ an act with range $f[X] = \{z_1, \dots, z_k\}$ such that for every $n \in \{1, \dots, k\}$ there is some measurable set E_n in Y such that $f^{-1}[\{z_n\}] = \varphi^{-1}[E_n]$. Then there is an act $f' \in FY$ such that $f = f' \circ \varphi$.

Proof. We use an induction to define $H_1 = E_1$ and $H_{n+1} = E_{n+1} \setminus \bigcup_{m=1}^n H_m$ for $n \in \{1, \dots, k-1\}$. We also define $H_0 = Y \setminus \bigcup_{m=1}^k H_m$. It is clear that all these sets H_0, H_1, \dots, H_k are measurable in Y . Moreover, it is also obvious from the definition that they partition the set Y . Let $z_0 \in Z$ be an arbitrary element of the non-empty set of outcomes. Define the act $f' \in FY$ such that $f'(y) = z_n$ if $y \in H_n$. Because the sets H_0, H_1, \dots, H_k partition Y this is well-defined as a function. It is a measurable function because the sets H_0, H_1, \dots, H_k are measurable. To see that $f = f' \circ \varphi$, fix any $x \in X$ and let n be such that $f(x) = z_n$. Because $f^{-1}[\{z_n\}] = \varphi^{-1}[E_n]$ we then have that $x \in \varphi^{-1}[E_n]$. Also, for every $m \neq n$ the sets $f^{-1}[\{z_n\}]$ and $f^{-1}[\{z_m\}] = \varphi^{-1}[E_m]$ are disjoint and hence $x \notin \varphi^{-1}[E_m]$. It follows that $\varphi(x) \in H_n$ and thus $f'(\varphi(x)) = z_n = f(x)$. \square

The next lemma states a property that also plays an important role in Proposition 1 of [DT08]:

Lemma 5. Consider a countable cochain $(X_n, \xi_n)_{n \in \omega}$, with coherence morphisms $\xi_n : X_{n+1} \rightarrow X_n$, and let X_ω , with projections $\zeta_n : X_\omega \rightarrow X_n$, be its limit. Then, for every act $f \in FX_\omega$ there is a natural number $n \in \omega$ such that for every natural number $m \geq n$ there is an act $f' \in FX_m$ such that $f = f' \circ \zeta_m$.

Proof. Because $f : X_\omega \rightarrow Z$ is a finite step function we can enumerate its range as $f[X] = \{z_1, \dots, z_k\}$. Then observe that for every $i \in \{1, \dots, k\}$ the set $f^{-1}[\{z_i\}]$ is measurable in X_ω , since f is a measurable and Z has the discrete algebra. By the definition of the algebra on X_ω this means that for every $i \in \{1, \dots, k\}$ there is some $n_i \in \omega$ such that $f^{-1}[\{z_i\}] = \zeta_{n_i}^{-1}[E_i]$ for some measurable set E_i in X_{n_i} . Let n be the maximum of all n_i for $i \in \{1, \dots, k\}$.

Consider then any $m \geq n$ and for all $i \in \{1, \dots, k\}$ define $E_i^m = (\xi_{n_i}^m)^{-1}[E_i] \subseteq X_m$. These E_i^m are measurable in X_m because $\xi_{n_i}^m$ is a measurable function. By the definition of the limit of a cochain we also have that $\zeta_{n_i} = \xi_{n_i}^m \circ \zeta_m$ for every $i \in \{1, \dots, k\}$. Hence $f^{-1}[\{z_i\}] = \zeta_{n_i}^{-1}[E_i] = (\xi_{n_i}^m \circ \zeta_m)^{-1}[E_i] = \zeta_m^{-1}[(\xi_{n_i}^m)^{-1}[E_i]] = \zeta_m^{-1}[E_i^m]$. It thus follows with Lemma 4 that there is an act $f' \in X_m$ such that $f = f' \circ \zeta_m$. \square

Proposition 4. F preserves limits of countable epic cochains. That is, it maps limits of cochains of uncertainty spaces to colimits of chains of sets.

Proof. Consider a cochain $(X_n, \xi_n)_{n \in \omega}$ of uncertainty spaces and let X_ω , together with projections $\zeta_n : X_\omega \rightarrow X_n$, be its limit. We show that FX_ω , together with the inclusions $F\zeta_n : FX_n \rightarrow FX_\omega$ has the universal property of the colimit of the chain $(FX_n, F\xi_n)_{n \in \omega}$. For this purpose consider any set T together with functions $h_n : X_n \rightarrow T$ such that $h_n = h_{n+1} \circ F\xi_n$ for each $n \in \omega$. We need to define the function $u : FX_\omega \rightarrow T$ such that $h_n = u \circ F\zeta_n$ for all $n \in \omega$ and show that it is unique with that property.

To define $u(f)$ for some act $f \in FX_\omega$ we use Lemma 5. From this lemma it follows that there is some k and act $f' \in FX_k$ such that $f = f' \circ \zeta_k = F\zeta_k(f')$. We then set $u(f) = h_k(f')$.

To check that this satisfies $h_n(g) = u \circ F\zeta_n(g)$ for all $n \in \omega$ and $g \in FX_n$ consider the act $F\zeta_n(g) \in FX_\omega$. By definition of u we have that $u(F\zeta_n(g)) = h_k(f')$ for some $f' \in FX_k$ such that $F\zeta_n(g) = F\zeta_k(f')$. We need to show that $h_k(f') = h_n(g)$. Assume that $n \leq k$. This is without loss of generality because we are only using the completely symmetric fact that $F\zeta_n(g) = F\zeta_k(f')$. Because $\zeta_n = \xi_n^k \circ \zeta_k$, it follows from $F\zeta_n(g) = F\zeta_k(f')$ that $f' \circ \zeta_k = g \circ \zeta_n = g \circ \xi_n^k \circ \zeta_k$. We obtain that $f' = g \circ \xi_n^k = F\xi_n^k(g)$ because ζ_k is epic. From this we can then conclude that $h_k(f') = h_n(g)$ since $h_k = h_n \circ F\xi_n^k$.

That u is unique also follows from Lemma 5. The possible values of $u(f)$ are completely determined because $f = F\zeta_n(f')$ for some n and $f' \in X_n$ and we need to ensure that $h_n(f') = u \circ F\zeta_n(f')$. \square

Lemma 6. Consider a chain $(X_n, \xi_n)_{n \in \omega}$ of sets and let X_ω , with inclusions $\iota_n : X_n \rightarrow X_\omega$ for all n , be its colimit. Then for each finite $K \subseteq X_\omega$ there is an $m \in \omega$ and a $K' \subseteq X_m$ such that $\iota_m[K'] = K$.

Proof. By the definition of the colimit of a countable chain we have that for every $k \in X_\omega$ there exists some $k' \in X_{m_k}$ such that $\iota_{m_k}(k') = k$. Let m be the maximum of the finitely many m_k for all $k \in K$ and let then $K' = \{\xi_m^{m_k}(k') \in X_m \mid k \in K\}$. It then holds that $\iota_m[K'] = K$ because $\iota_m \circ \xi_m^{m_k}(k') = \iota_{m_k}(k') = k$ for every $k \in K$. \square

Theorem 5. \mathcal{C} preserves colimits of countable chains. That is, it maps colimits of chains of sets onto limits of cochains of uncertainty spaces.

Proof. Consider a chain $(X_n, \xi_n)_{n \in \omega}$ of sets and let X_ω , with inclusions $\iota_n : X_n \rightarrow X_\omega$ for all n , be its colimit. We show that $\mathcal{C}X_\omega$, with the $\mathcal{C}\iota_n : \mathcal{C}X_\omega \rightarrow \mathcal{C}X_n$ as the projections, has the universal property of the limit of the cochain $(\mathcal{C}X_n, \mathcal{C}\xi_n)_{n \in \omega}$. Hence, consider any uncertainty space T together with measurable functions $\varphi_n : T \rightarrow \mathcal{C}X_n$, satisfying $\varphi_n = \mathcal{C}\xi_n \circ \varphi_{n+1}$, for all $n \in \omega$. We need to show that there is a unique measurable function $\mu : T \rightarrow \mathcal{C}X_\omega$ such that $\mathcal{C}\iota_n \circ \mu = \varphi_n$ for all $n \in \omega$.

We first describe how to define the choice function $\mu(t) \in \mathcal{C}X_\omega$ on a finite set $K \subseteq X_\omega$. Let $m \in \omega$ be the least number such that there exists a $K' \subseteq X_m$ such that $\iota_m[K'] = K$. Such a number exists because of Lemma 6. Also observe that the choice of m and K' only depends on K but not on t . This is a property which we exploit below to show that μ is measurable. We then set the value of the choice function $\mu(t)$ on K to be

$$\mu(t)(K) = \iota_m[\varphi_m(t)(K')], \quad (2)$$

where $\varphi_m(t)(K') \subseteq K'$ denotes the elements of K' selected by the choice function $\varphi_m(t) \in \mathcal{C}X_m$. This is well-defined because we have $\iota_m[\varphi_m(t)(K')] \subseteq \iota_m[K'] = K$.

To see that μ is measurable it suffices to check that the preimage $\mu^{-1}[B_L^K] \subseteq T$ of a basic measurable set $B_L^K = \{C \in \mathcal{C}X_\omega \mid C(K) \subseteq L\}$ is measurable in T . Hence, fix finite K and L with $L \subseteq K$ and let $m \in \omega$ and K' be defined from K as explained above. We are now going to show that $\mu^{-1}[B_L^K] = \varphi_m^{-1}[B_{L'}^{K'}]$, where $L' = \iota_m^{-1}[L]$. Because $\varphi_m : T \rightarrow \mathcal{C}X_m$ is assumed to be measurable, it then follows that $\mu^{-1}[B_L^K]$ is measurable too. By unfolding the definitions one sees that the claim that $\mu^{-1}[B_L^K] = \varphi_m^{-1}[B_{L'}^{K'}]$ is equivalent to the claim that for all $t \in T$

$$\mu(t)(K) \subseteq L \quad \text{iff} \quad \varphi_m(t)(K') \subseteq \iota_m^{-1}[L].$$

By the definition of $\mu(t)$ above we see that the left side of this equivalence is the same as $\iota_m[\varphi_m(t)(K')] \subseteq L$, which is clearly equivalent to $\varphi_m(t)(K') \subseteq \iota_m^{-1}[L]$.

We need to verify that $\varphi_n = \mathcal{C}\iota_n \circ \mu$ for all $n \in \omega$. This requires that for every $t \in T$ the choice functions $\varphi_n(t)$ and $\mathcal{C}\iota_n \circ \mu(t)$ are the same. Hence, they need to have the same value on every finite set $L' \subseteq X_n$.

First let $K = \iota_n[L']$ and recall the definitions to see that

$$(\mathcal{C}\iota_n \circ \mu(t))(L') = \iota_n^{-1}[\mu(t)(\iota_n[L'])] \cap L' = \iota_n^{-1}[\iota_m[\varphi_m(t)(K')]] \cap L', \quad (3)$$

for the $m \in \omega$ and the finite set $K' \subseteq X_m$ that are defined from K as described above. Our choice of K' above ensures that $\iota_m[K'] = K = \iota_n[L']$. Hence for every $k' \in K'$ there is some $l' \in L'$ such that $\iota_m(k') = \iota_n(l')$ and conversely for every $l' \in L'$ there is a $k' \in K'$ such that $\iota_m(k') = \iota_n(l')$. By the definition of identity for elements in the colimit X_ω it follows that for each such pair $k' \in X_m$ and $l' \in X_n$ with $\iota_m(k') = \iota_n(l')$ there is some $j \in \omega$ such that $\xi_j^m(k') = \xi_j^n(l')$. Let i be the maximum of all those $j \in \omega$, which exists because there are finitely many pairs $(k', l') \in K' \times L'$. It then clearly holds that $\xi^m(k') = \xi^n(l')$, whenever $\iota_m(k') = \iota_n(l')$ for $(k', l') \in K' \times L'$. Define then $L \subseteq X$ to be the finite set

$$L = \xi^m[K'] = \xi^n[L'].$$

Because $\varphi_n = \mathcal{C}\xi^n \circ \varphi$ we obtain that

$$\varphi_n(t)(L') = (\xi^n)^{-1}[\varphi(t)(\xi_i^n[L'])] \cap L' = (\xi^n)^{-1}[\varphi(t)(L)] \cap L'.$$

Similarly, because $\varphi_m = \mathcal{C}\xi^m \circ \varphi$ we obtain

$$\varphi_m(t)(K') = (\xi^m)^{-1}[\varphi_m(t)(\xi^m[K'])] \cap K' = (\xi^m)^{-1}[\varphi(t)(L)] \cap K'.$$

To prove $\varphi_n(t)(L') = (\mathcal{C}\iota_n \circ \mu(t))(L')$ it thus suffices by (3) to show that

$$(\xi^n)^{-1}[U] \cap L' = \iota_n^{-1}[\iota_m((\xi^m)^{-1}[U] \cap K')] \cap L' \quad (4)$$

for the set $U = \varphi(t)(L) \subseteq L$.

For the left-to-right inclusion of (4) consider any $l' \in L'$ such that $\xi^n(l') \in U$. We need that $l' \in \iota_n^{-1}[\iota_m((\xi^m)^{-1}[U] \cap K')]$. From the definition of L we get that there is then some $k' \in K'$ such that $\xi^m(k') = \xi^n(l')$. Hence $k' \in (\xi^m)^{-1}[U] \cap K'$ and $\iota_n(l') = \iota_m(k')$. The latter two directly entail that $l' \in \iota_n^{-1}[\iota_m((\xi^m)^{-1}[U] \cap K')]$.

For the right-to-left inclusion of (4) consider any $l' \in L'$ such that $\iota_n(l') = \iota_m(k')$ for some $k' \in (\xi^m)^{-1}[U] \cap K'$. With the definition of i above it then follows from $\iota_n(l') = \iota_m(k')$ that $\xi^n(l') = \xi^m(k') \in U$. Hence, $l' \in (\xi^n)^{-1}[U]$ and we are done.

Lastly, we show that $\mu(t)$ is completely determined by the requirement that $\varphi_n(t) = \mathcal{C}\iota_n \circ \mu(t)$ for all $n \in \omega$. Consider any $\nu : T \rightarrow \mathcal{C}X_\omega$ such that $\varphi_n(t) = \mathcal{C}\iota_n \circ \nu(t)$ for all $n \in \omega$. We argue that then $\nu(t)(K) = \mu(t)(K)$ for the μ defined as above and all finite $K \subseteq X_\omega$. Let $m \in \omega$ and $K' \subseteq X_m$ as in the definition of μ above, i.e., such that $\iota_m[K'] = K$. The requirement that $\varphi_m(t) = \mathcal{C}\iota_m \circ \nu(t)$ means that

$$\varphi_m(t)(K') = \iota_m^{-1}[\nu(t)(\iota_m[K'])] \cap K' = \iota_m^{-1}[\nu(t)(K)] \cap K'.$$

One can see that this entails that

$$\iota_m[\varphi_m(t)(K')] = \nu(t)(K).$$

The left-to-right inclusion is trivial and the other inclusion follows because $\nu(t)(K) \subseteq K$ and $\iota_m[K'] = K$. Therefore, we have shown that $\nu(t)(K)$ equals the expression that we use in (2) to define $\mu(t)(K)$. \square

As an immediate consequence of the previous two propositions we obtain the main result of this section.

D Proofs for Section 3

Theorems 1 and 2 follow from the results about Γ together with well-known results [Wor99] about the terminal coalgebra. In the following, we sketch how these results from the general theory of coalgebras apply to the setting of choice structures.

D.1 Preliminary observations

Theorems 1 and 2 concern the universal choice structure that is obtained from the choice hierarchies as described in Section 3.2. Let us first see how this fits into the coalgebraic set-up.

One can view the two uncertainty spaces $\Omega_{i,n}$ and $\Omega_{j,n}$ on the n -th level in the choice hierarchy as defining an object $\Omega_n = (\Omega_{i,n}, \Omega_{j,n})$ in the category \mathbf{Unc}^2 . In fact one can define these objects directly, just using the functor Γ^\heartsuit as follows: We start with $\Omega_0 = \top$, where \top is the terminal object in \mathbf{Unc}^2 , and then define inductively $\Omega_{n+1} = \Gamma^\heartsuit\Omega_n$. It is easy to see that with the exception of the 0-th level, which is omitted from the discussion in the main text, this yields the same sequence of pairs of uncertainty spaces as defined in Section 3.2. Similarly, the coherence morphism $\xi_{i,n}$ and $\xi_{j,n}$ can also be defined directly in \mathbf{Unc}^2 with $\xi_0 = !_{\Gamma^\heartsuit\top} : \Gamma^\heartsuit\top \rightarrow \top$ and inductively $\xi_{n+1} = \Gamma^\heartsuit\xi_n : \Gamma^\heartsuit\Omega_{n+1} \rightarrow \Gamma^\heartsuit\Omega_n$. Thus, the sequence $(\Omega_n, \xi_n)_{n \in \omega}$ forms a countable cochain in \mathbf{Unc}^2 .

The definition of the uncertainty spaces Ω_i and Ω_j of types in the universal choice structures from Section 3.2 is such that these are precisely the limits of the countable cochains $(\Omega_{i,n}, \xi_{i,n})_{n \in \omega}$ and $(\Omega_{j,n}, \xi_{j,n})_{n \in \omega}$ in \mathbf{Unc} . Because limits in \mathbf{Unc}^2 are computed component-wise it follows that the object $\Omega = (\Omega_i, \Omega_j)$ in \mathbf{Unc}^2 is the limit of the countable cochain $(\Omega_n, \xi_n)_{n \in \omega}$. Also recall that the limit comes with projections $\zeta_n = (\zeta_{i,n}, \zeta_{j,n}) : \Omega \rightarrow \Omega_n$ back into the chain such that $\zeta_n = \xi_n \circ \zeta_{n+1}$.

The crucial observation behind Theorems 1 and 2 is then that the functor Γ^\heartsuit on \mathbf{Unc}^2 preserves the limit of the epic cochain (Ω_n, ξ_n) . This follows from Corollary 2, which states that Γ^\heartsuit preserves limits of countable epic cochains. To apply Corollary 2 to the cochain $(\Omega_n, \xi_n)_{n \in \omega}$ we need to check that the coherence morphisms ξ_n for all $n \in \omega$ are epic. This can be checked directly by an induction over the definition of the ξ_n . In the base step $\xi_0 = (\xi_{i,0}, \xi_{j,0})$, and $\xi_{i,0} = \Gamma\pi_1$, where π_1 is the projection out of a product. Because this projection is epic, and by the argument in Section C.1 the functor Γ preserves epic morphisms, it follows that $\xi_{i,0}$ is epic. We reason analogously for $\xi_{j,0}$. That $\xi_{n+1} = (\xi_{i,n+1}, \xi_{j,n+1}) = (\Gamma(\text{id}_{A_j} \times \xi_{j,n}), \Gamma(\text{id}_{A_i} \times \xi_{i,n}))$ is epic, also follows easily because Γ preserves epic morphisms.

D.2 Theorem 1

Next consider the object $\Gamma^\heartsuit\Omega$ of \mathbf{Unc}^2 . For this object we can define morphisms into the cochain $(\Omega_n, \xi_n)_{n \in \omega}$ by setting $\tau_0 = !_{\Gamma^\heartsuit\Omega} : \Gamma^\heartsuit\Omega \rightarrow \Omega_0$ and $\tau_{n+1} = \Gamma^\heartsuit\zeta_n : \Gamma^\heartsuit\Omega \rightarrow \Gamma^\heartsuit\Omega_n$, which satisfy $\tau_n = \xi_n \circ \tau_{n+1}$ for all n . From Theorem 4 it follows that $\Gamma^\heartsuit\Omega$ together with the projections $\Gamma^\heartsuit\zeta_n = \tau_{n+1}$ is a colimit of the cochain $(\Gamma^\heartsuit\Omega_n, \Gamma^\heartsuit\xi_n)_{n \in \omega}$, which by definition is the same as the chain $(\Omega_{n+1}, \xi_{n+1})_{n \in \omega}$. We can use this observation to show that $\Gamma^\heartsuit\Omega$ together with the τ_n also satisfies the universal property of the limit of the cochain $(\Omega_n, \xi_n)_{n \in \omega}$. To this aim take any further object T of \mathbf{Unc}^2 together with morphisms $g_n : T \rightarrow \Omega_n$ such that $g_n = \xi_n \circ g_{n+1}$ for all $n \in \omega$. We now just consider this g_n without g_0 as morphism into the chain $(\Gamma^\heartsuit\Omega_n, \Gamma^\heartsuit\xi_n)_{n \in \omega}$ satisfying that $g_{n+1} = \xi_{n+1} \circ g_{n+2} = \Gamma^\heartsuit\xi_n \circ g_{n+1}$ for all $n \in \omega$. Because $\Gamma^\heartsuit\Omega$ is a limit of this chain there must be then a unique morphism $u : T \rightarrow \Gamma^\heartsuit\Omega$ such that $g_{n+1} = \Gamma^\heartsuit\zeta_n \circ u = \tau_{n+1} \circ u$ for all $n \in \omega$. Additionally, we have that $g_0 = \tau_0 \circ u$ holds trivially because on both sides of the equation we have a morphism into the terminal object $\top = \Omega_0$ of \mathbf{Unc}^2 .

We have now seen that both Ω and $\Gamma^\heartsuit\Omega$ satisfy the universal property of limit for the cochain $(\Omega_n, \xi_n)_{n \in \omega}$. Theorem 1 then follows because there can be only one such object up to isomorphism.

More precisely, the isomorphism is given by the unique morphism $\mu = (\mu_i, \mu_j) : \Omega \rightarrow \Gamma^\heartsuit\Omega$ such that $\zeta_n = \tau_n \circ \mu$ for all $n \in \omega$, which exists because Ω is a limit. If one considers the components of this morphism one obtains the isomorphism $\mu_i : \Omega_i \rightarrow \Gamma(A_j \times \Omega_j)$ and $\mu_j : \Omega_j \rightarrow \Gamma(A_i \times \Omega_i)$ that are referred to in the formulation of Theorem 1.

Remark 2. One can generate a topology on Ω_i , using the algebra of events as the basis. Because of our assumption that Z , A_i and A_j are finite it follows that all the $\Omega_{i,n}$ are finite. Thus, Ω_i is what is called a filtered limit of finite spaces, which is well-known to be a Stone space, meaning that it is compact, Hausdorff and has a basis of clopen sets. One can also show also that this topology has a countable basis, because for every of the countably many spaces $\Omega_{i,n}$ there are only finitely

many events in Ω_i that are preimages of sets in $\Omega_{i,n}$. Compact Hausdorff spaces with a countable basis are called simple spaces in [DT08].

Remark 3. For the result of this section, and the universality of the universal choice structure proved in the next section, it is not needed that the uncertainty spaces Z , A_i , and A_j are finite. The preservation of limits of epic cochains, which is at the core of these results, holds for arbitrary uncertainty spaces. However, for our proof of Lemma 4 we need that acts are defined as finite step functions and that Z carries the discrete algebra. The finiteness of Z , A_i and A_j is required for the results about non-redundancy from Section 3.4 and the embedding of the universal preference structure from Section 4.3.

D.3 Theorem 2

We now prove Theorem 2. Consider an arbitrary choice structure $\mathcal{X} = (T_i, T_j, \theta_i, \theta_j)$, presented as a coalgebra $(T, \theta) = ((T_i, T_j), (\theta_i, \theta_j))$ for Γ^\heartsuit . We need to show that there is a unique morphism v from (T, θ) to the universal choice structure (Ω, μ) .

To prove the existence of v first consider the measurable functions $v_0 = !_T : T \rightarrow \Omega_0$ and the inductively defined $v_{n+1} = \Gamma^\heartsuit v_n \circ \theta : T \rightarrow \Omega_{n+1}$. By an induction over n we can show that these measurable functions satisfy $v_n = \xi_n \circ v_{n+1}$. In the base case this is clear because there is only one morphism from T to the terminal object $\Omega_0 = \top$. For the inductive step we calculate as follows:

$$v_{n+1} = \Gamma^\heartsuit v_n \circ \theta = \Gamma^\heartsuit(\xi_n \circ v_{n+1}) = \Gamma^\heartsuit \xi_n \circ \Gamma^\heartsuit v_{n+1} \circ \theta = \xi_{n+1} \circ v_{n+2}.$$

Because Ω is defined as the limit of the sequence $(\Omega_n, \xi_n)_{n \in \omega}$ it follows that there is a unique morphism $v : T \rightarrow \Omega$ with the property that for all $n \in \omega$

$$v_n = \zeta_n \circ v. \quad (5)$$

It remains to be seen that v is a morphism of choice structures and that it is unique with this property.

To show that $v : T \rightarrow \Omega$ is a morphism from (T, θ) to (Ω, μ) we need to verify that $\mu \circ v = \Gamma^\heartsuit v \circ \theta$. We show this by proving that $v_n = \tau_n \circ \mu \circ v$ and $v_n = \tau_n \circ \Gamma^\heartsuit v \circ \theta$ hold for all $n \in \omega$. The equality $\mu \circ v = \Gamma^\heartsuit v \circ \theta$ then follows from the universal property for the limit $\Gamma^\heartsuit \Omega$ of $(\Omega_n, \xi_n)_{n \in \omega}$, with projections $\tau_n : \Gamma^\heartsuit \Omega \rightarrow \Omega_n$.

To see that $v_n = \tau_n \circ \mu \circ v$ we use that by its definition in Section D.2 μ satisfies $\tau_n \circ \mu = \zeta_n$. The claim then reduces to (5).

To see that $v_n = \tau_n \circ \Gamma^\heartsuit v \circ \theta$ we distinguish two cases. If $n = 0$ we have that both v_0 and $\tau_0 \circ \Gamma^\heartsuit v \circ \theta$ are morphisms from T to the terminal object $\Omega_0 = \top$ of \mathbf{Unc}^2 . By the universal property of the terminal object they must be equal.

If n is strictly positive we can unfold the definition of v_{n+1} , use (5) and then unfold the definition of τ_{n+1} to obtain the following computation

$$v_{n+1} = \Gamma^\heartsuit v_n \circ \theta = \Gamma^\heartsuit(\zeta_n \circ v) \circ \theta = \Gamma^\heartsuit \zeta_n \circ \Gamma^\heartsuit v \circ \theta = \tau_{n+1} \circ \Gamma^\heartsuit v \circ \theta.$$

To prove that v is the only morphism of choice structures from (T, θ) to (Ω, μ) consider any other morphism $v' : T \rightarrow \Omega$ in \mathbf{Unc}^2 such that $\mu \circ v' = \Gamma^\heartsuit v' \circ \theta$. We show that then $v_n = \zeta_n \circ v'$ for all $n \in \omega$, from which it follows that $v' = v$ because v is defined as the unique morphism with the property (5). To prove $v_n = \zeta_n \circ v'$ we use an induction on n . In the base case we again have that v_0 and $\zeta_0 \circ v'$ must be equal because they are both morphism from T to the terminal object $\Omega_0 = \top$ of \mathbf{Unc}^2 . In the inductive step we use the following computation:

$$\begin{aligned} v_{n+1} &= \Gamma^\heartsuit v_n \circ \theta && \text{definition of } v_{n+1} \\ &= \Gamma^\heartsuit(\zeta_n \circ v') \circ \theta && \text{induction hypothesis} \\ &= \Gamma^\heartsuit \zeta_n \circ \Gamma^\heartsuit v' \circ \theta && \Gamma^\heartsuit \text{ functor} \\ &= \tau_{n+1} \circ \Gamma^\heartsuit v' \circ \theta && \text{definition of } \tau_{n+1} \\ &= \tau_{n+1} \circ \mu \circ v' && \text{assumption on } v' \\ &= \zeta_{n+1} \circ v' && \text{uniqueness property of } \mu \end{aligned}$$

D.4 Results for Section 3.4

In this part of the appendix we prove our results about non-redundancy.

D.4.1 Well-definedness of Definition 3

To explain the definition of the algebra of observable events more extensively, fix the choice structure $\mathcal{X} = (T_i, T_j, \theta_i, \theta_j)$ and consider pairs of algebras $(\mathcal{B}_i, \mathcal{B}_j)$ such that \mathcal{B}_i is an algebra over the set T_i and \mathcal{B}_j is an algebra over the set T_j . Define an order \leq over such pairs of algebras such that $(\mathcal{B}_i, \mathcal{B}_j) \leq (\mathcal{B}'_i, \mathcal{B}'_j)$ iff $\mathcal{B}_i \subseteq \mathcal{B}'_i$ and $\mathcal{B}_j \subseteq \mathcal{B}'_j$. Definition 3 then defines $(\mathcal{B}_{\mathcal{X},i}^{ob}, \mathcal{B}_{\mathcal{X},j}^{ob})$ as the least element $(\mathcal{B}_i, \mathcal{B}_j)$ in this order \leq which satisfies that

1. $B_{\mathcal{X},i}^{ob}(K, F) \in \mathcal{B}_i$ for all finite $K, L \subseteq F(A_j \times (T_j, \mathcal{B}_j))$, and
2. $B_{\mathcal{X},j}^{ob}(K, F) \in \mathcal{B}_j$ for all finite $K, L \subseteq F(A_i \times (T_i, \mathcal{B}_i))$.

Let O be the set of all pairs of algebras that satisfy these two conditions. The well-definedness of $(\mathcal{B}_{\mathcal{X},i}^{ob}, \mathcal{B}_{\mathcal{X},j}^{ob})$ now amounts to the following claim:

Proposition 5. O contains a (unique) least element.

Proof. Define the pair of families of sets $(\mathcal{L}_i, \mathcal{L}_j)$ where $\mathcal{L}_i = \bigcap \{ \mathcal{B}_i \mid (\mathcal{B}_i, \mathcal{B}_j) \in O \}$ and $\mathcal{L}_j = \bigcap \{ \mathcal{B}_j \mid (\mathcal{B}_i, \mathcal{B}_j) \in O \}$. It is relatively easy to see that $(\mathcal{L}_i, \mathcal{L}_j) \leq (\mathcal{B}_i, \mathcal{B}_j)$ for all $(\mathcal{B}_i, \mathcal{B}_j) \in O$. We argue that $(\mathcal{L}_i, \mathcal{L}_j)$ is again a pair of algebras and it is in O . Thus, $(\mathcal{L}_i, \mathcal{L}_j)$ is the least element of O . We leave it to the reader to check that $(\mathcal{L}_i, \mathcal{L}_j)$ is again a pair of algebras. This essentially reduced to the well-known fact that the arbitrary intersection of algebras is again an algebra.

To show that $(\mathcal{L}_i, \mathcal{L}_j) \in O$ we need to show that it satisfies the two conditions of the list above. We only check the first one that $B_{\mathcal{X},i}^{ob}(K, F) \in \mathcal{L}_i$ for all finite $K, L \subseteq F(A_j \times (T_j, \mathcal{L}_j))$. To this aim fix any finite $K, L \subseteq F(A_j \times (T_j, \mathcal{L}_j))$. To show that $B_{\mathcal{X},i}^{ob}(K, F) \in \mathcal{L}_i$ we need to show that $B_{\mathcal{X},i}^{ob}(K, F) \in \mathcal{B}_i$ for all $(\mathcal{B}_i, \mathcal{B}_j) \in O$. This would follow because $(\mathcal{B}_i, \mathcal{B}_j)$ also satisfies the first item in the list above, provided that we can show that $K, L \subseteq F(A_j \times (T_j, \mathcal{B}_j))$. To obtain the latter we argue that $F(A_j \times (T_j, \mathcal{L}_j)) \subseteq F(A_j \times (T_j, \mathcal{B}_j))$. Because $\mathcal{L}_j \subseteq \mathcal{B}_j$ it follows that every measurable set in $A_j \times (T_j, \mathcal{L}_j)$ is also measurable in $A_j \times (T_j, \mathcal{B}_j)$. Hence, a function that is measurable for the domain $A_j \times (T_j, \mathcal{L}_j)$ is also measurable for the domain $A_j \times (T_j, \mathcal{B}_j)$. \square

D.4.2 Proof of Proposition 1

To prove Proposition 1 we need to perform another inductive argument over the definition of the terminal sequence from which the universal choice structure \mathcal{U} was defined. Let us first recall relevant facts from Sections D.1 and D.3 that will be used in the following arguments. For every n we had an approximant Ω_n in \mathbf{Unc}^2 such that Ω_0 is the terminal object of \mathbf{Unc}^2 and $\Omega_{n+1} = \Gamma^\heartsuit \Omega_n$. Concretely, this meant that $\Omega_{i,0} = \top$ is the one element uncertainty space and $\Omega_{i,n+1} = \Gamma(A_j \times \Omega_{j,n})$, and similarly with i and j swapped. Given any choice structure $\mathcal{X} = (T, \theta)$, where $T = (T_i, T_j)$ and $\theta = (\theta_i, \theta_j)$ are in \mathbf{Unc}^2 , we then defined the maps $v_n : T \rightarrow \Omega_n$ such that $v_0 = !_T$ is the unique morphism into the terminal object Ω_0 of \mathbf{Unc}^2 and $v_{n+1} = \Gamma^\heartsuit v_n \circ \theta$. If we spell out the latter in terms of measurable functions we have that

$$v_{i,n+1} = \Gamma(\text{id}_{A_j} \times v_{j,n}) \circ \theta_i \quad \text{and} \quad v_{j,n+1} = \Gamma(\text{id}_{A_i} \times v_{i,n}) \circ \theta_j. \quad (6)$$

We then used the properties of the limit Ω to obtain a the morphism $v : T \rightarrow \Omega$, which is unique for the property that $v_n = \zeta_n \circ v$. Here, $\zeta_n : \Omega \rightarrow \Omega_n$ is the projection from the limit into the cochain. For measurable functions this property means that

$$v_{i,n} = \zeta_{i,n} \circ v_i \quad \text{and} \quad v_{j,n} = \zeta_{j,n} \circ v_j. \quad (7)$$

Lemma 7. Let $\mathcal{X} = (T_i, T_j, \theta_i, \theta_j)$ be any choice structure and for every $n \in \omega$ consider $v_{i,n} : T_i \rightarrow \Omega_{i,n}$ as defined above. Then it holds for every set $A \subseteq \Omega_{i,n}$ that $v_{i,n}^{-1}[A] \in \mathcal{B}_{\mathcal{X},i}^{ob}$. The same holds with j in place of i .

Proof. The proof is an induction on n . In the base case we have that $\Omega_{i,0}$ is the one element uncertainty space. Thus, $v_{i,0}^{-1}[A]$ is either the empty set or the whole set of all types in T_i . Thus, it is in the algebra $\mathcal{B}_{\mathcal{X},i}^{ob}$.

To prove the inductive step we show below that $v_{i,n+1}^{-1}[B_L^K] \in \mathcal{B}_{\mathcal{X},i}^{ob}$ holds for for all $K, L \subseteq F(A_j \times \Omega_{j,n})$, where $B_L^K = \{C \in \Omega_{j,n+1} \mid C(K) \subseteq L\}$. Let us see first how from this it then follows that $v_{i,n+1}^{-1}[A] \in \mathcal{B}_{\mathcal{X},i}^{ob}$ holds for every arbitrary set $A \subseteq \Omega_{n+1,i}$: Because $\Omega_{i,n+1}$ is finite and any two distinct choice functions in $\Omega_{i,n+1} = \Gamma(A_j \times \Omega_{j,n})$ can be separated by some such set B_L^K , it follows from a standard argument that every set $A \subseteq \Omega_{n+1,i}$ can be written as a finite union of intersections of sets of the form B_L^K and complements of such sets. Because these unions, intersections and complements are preserved under taking preimages and $\mathcal{B}_{\mathcal{X},i}^{ob}$ is an algebra it follows that $v_{i,n}^{-1}[A] \in \mathcal{B}_{\mathcal{X},i}^{ob}$.

We now prove that $v_{i,n+1}^{-1}[B_L^K] \in \mathcal{B}_{\mathcal{X},i}^{ob}$ holds for for all $K, L \subseteq F(A_j \times \Omega_{j,n})$. First, observe that by the induction hypothesis we have that if $f \in F(A_j \times \Omega_{j,n})$ then $f \circ (\text{id}_{A_j} \times v_{j,n}) \in F(A_j \times (\Omega_j, \mathcal{B}_{\mathcal{X},j}^{ob}))$. This holds because by the inductive hypothesis the preimage of any set at $\Omega_{j,n}$ under $v_{j,n}$ is measurable in the uncertainty space $(T_j, \mathcal{B}_{\mathcal{X},j}^{ob})$. Thus, the function $f \circ (\text{id}_{A_j} \times v_{j,n})$ is measurable even when we take its domain to be $A_j \times (T_j, \mathcal{B}_{\mathcal{X},j}^{ob})$. Then, fix the sets $K, L \subseteq F(A_j \times \Omega_{j,n})$ and define the finite sets

$$K' = \{f \circ (\text{id}_{A_j} \times v_{j,n}) \mid f \in K\} \quad \text{and} \quad L' = \{f \circ (\text{id}_{A_j} \times v_{j,n}) \mid f \in L\}.$$

From the observation in the beginning of the paragraph it follows that that $K', L' \subseteq F(A_j \times (T_j, \mathcal{B}_{\mathcal{X},j}^{ob}))$.

To show that $v_{i,n+1}^{-1}[B_L^K] \in \mathcal{B}_{\mathcal{X},i}^{ob}$, it suffices by the definition of $\mathcal{B}_{\mathcal{X},i}^{ob}$ to show that $v_{i,n+1}^{-1}[B_L^K] = B_{\mathcal{X},i}^{ob}(K', L')$. Unfolding the definitions shows that this amounts to the claim that for all $t \in T_i$ we have $v_{i,n+1}(t)(K) \subseteq L$ iff $\theta_i(t)(K') \subseteq L'$. To see that this is the case consider the following chain of equivalences:

$$\begin{aligned} v_{i,n+1}(t)(K) \subseteq L &\text{ iff } (\Gamma(\text{id}_{A_j} \times v_{j,n}) \circ \theta_i(t))(K) \subseteq L && \text{by (6)} \\ &\text{ iff } \{f \in K \mid f \circ (\text{id}_{A_j} \times v_{j,n}) \in \theta_i(t)(K')\} \subseteq L && \text{definition of } \Gamma \\ &\text{ iff } \theta_i(t)(K') \subseteq L' \end{aligned}$$

The last equivalence in this chain can be checked easily using the definitions of K' and L' . \square

Proposition 6. Let $\mathcal{X} = (T_i, T_j, \theta_i, \theta_j)$ be any choice structure and $v = (v_i, v_j)$ the unique morphism into the terminal choice structure $\mathcal{U} = (\Omega_i, \Omega_j, \mu_i, \mu_j)$. Let \mathcal{B}_i be the algebra of Ω_i . Then $\mathcal{B}_{\mathcal{X},i}^{ob} = \{v_i^{-1}[E] \mid E \in \mathcal{B}_i\}$. A similar claim holds with j for i .

Proof. Let $\mathcal{I}_i = \{v_i^{-1}[E] \mid E \in \mathcal{B}_i\}$ and $\mathcal{I}_j = \{v_j^{-1}[E] \mid E \in \mathcal{B}_j\}$. Because taking preimages preserves intersections and complements it follows that \mathcal{I}_i and \mathcal{I}_j are algebras. Before we argue that $\mathcal{B}_{\mathcal{X},i}^{ob} \subseteq \mathcal{I}_i$, define

$$\mathcal{J}_j = \{(\text{id}_{A_j} \times v_j)^{-1}[E] \mid E \text{ measurable in } A_j \times \Omega_j\}.$$

Observe that $A_j \times (T_j, \mathcal{I}_j) = (A_j \times T_j, \mathcal{J}_j)$. One can prove both inclusions of this equality by considering the generating cylinders in the product spaces $A_j \times (T_j, \mathcal{I}_j)$ and $A_j \times \Omega_j$. Then consider any act $f \in F(A_j \times (T_j, \mathcal{I}_j)) = F(A_j \times T_j, \mathcal{J}_j)$. We use Lemma 4 from Appendix C.3 to show that there is some act $f' \in F(A_j \times \Omega_j)$ such that

$$f = f' \circ (\text{id}_{A_j} \times v_j). \tag{8}$$

To see that all assumptions of Lemma 4 are satisfied, observe first that by the definition of \mathcal{J}_j the function $(\text{id}_{A_j} \times v_j)$ is measurable from $(A_j \times T_j, \mathcal{J}_j)$ to $A_j \times \Omega_j$. Moreover, because f is a measurable function from $(A_j \times T_j, \mathcal{J}_j)$ to the discrete uncertainty space Z we have that $f^{-1}[\{z\}]$ is measurable in $(A_j \times T_j, \mathcal{J}_j)$ for every $z \in Z$. By the definition of \mathcal{J}_j this means that for every $z \in Z$ there is some measurable set E_z in $A_j \times \Omega_j$ such that $f^{-1}[\{z\}] = (\text{id}_{A_j} \times v_j)^{-1}[E_z]$.

To prove that $\mathcal{B}_{\mathcal{X},i}^{ob} \subseteq \mathcal{I}_i$ we show that for all finite $K, L \subseteq F(A_j \times (T_j, \mathcal{I}_j))$ the set $B_{\mathcal{X},i}^{ob}(K, L)$ is in \mathcal{I}_i . This, and together with the analogous statement in which i and j are swapped, shows that \mathcal{I}_i and \mathcal{I}_j satisfy the properties from Definition 3. The claim that $\mathcal{B}_{\mathcal{X},i}^{ob} \subseteq \mathcal{I}_i$ then follows because $\mathcal{B}_{\mathcal{X},i}^{ob}$ and $\mathcal{B}_{\mathcal{X},j}^{ob}$ are defined to be the smallest algebras which satisfy these properties.

Consider any finite $K, L \subseteq F(A_j \times (T_j, \mathcal{I}_j))$. Define the finite sets $K', L' \subseteq F(A_j \times \Omega_j)$ such that $K' = \{f' \mid f \in K\}$ and $L' = \{f' \mid f \in L\}$, where f' is defined from f as in (8). It follows immediately from (8) that

$$K = \{g \circ (\text{id}_{A_j} \times v_j) \mid g \in K'\} \quad \text{and} \quad L = \{g \circ (\text{id}_{A_j} \times v_j) \mid g \in L'\}. \quad (9)$$

We are then going to show that

$$B_{\mathcal{X},i}^{ob}(K, L) = (\Gamma(\text{id}_{A_j} \times v_j) \circ \theta_i)^{-1}[B_{L'}^{K'}], \quad (10)$$

where $B_{L'}^{K'} = \{C \in \Gamma(A_j \times \Omega_j) \mid C(K') \subseteq L'\}$ is one of the basic measurable sets that generates the algebra on $\Gamma(A_j \times \Omega_j)$. Because v is a morphism of choice structures it follows that $B_{\mathcal{X},i}^{ob}(K, L) = (\mu_i \circ v_i)^{-1}[B_{L'}^{K'}] = v_i^{-1}[\mu_i^{-1}[B_{L'}^{K'}]]$. From this we obtain that $B_{\mathcal{X},i}^{ob}(K, L) \in \mathcal{I}_i = \{v_i^{-1}[E] \mid E \in \mathcal{B}_i\}$ because $\mu_i : \Omega_i \rightarrow \Gamma(A_j \times \Omega_j)$ is measurable and hence $\mu_i^{-1}[B_{L'}^{K'}] \in \mathcal{B}_i$.

To show (10) we unfold the definitions and see that it amounts to the claim that $\theta_i(t)(K) \subseteq L$ iff $(\Gamma(\text{id}_{A_j} \times v_j)(\theta_i(t)))(K') \subseteq L'$ holds for for all $t \in T_i$. Using the definition of Γ the right side of the equivalence in this claim unfolds to the statement

$$\{g \in K' \mid g \circ \Gamma(\text{id}_{A_j} \times v_j) \in \theta_i(t)(\{g \circ \Gamma(\text{id}_{A_j} \times v_j) \mid g \in K'\})\} \subseteq L'.$$

Using the definitions of K' and L' and the equalities from (9) one can easily see that this is equivalent to $\theta_i(t)(K) \subseteq L$.

To prove the other inclusion fix a measurable set $E \in \mathcal{B}_i$. We need to show that $v_i^{-1}[E] \in \mathcal{B}_{\mathcal{X},i}^{ob}$. By the definition of the algebra \mathcal{B}_i on the limit Ω_i it follows that there is some n such that $E = \zeta_{i,n}^{-1}[A]$ for some $A \subseteq \Omega_{i,n}$. By Lemma 7 we know that $v_{i,n}^{-1}[A] \in \mathcal{B}_{\mathcal{X},i}^{ob}$. From (7) we have that $v_{i,n} = \zeta_{i,n} \circ v_i$ and thus

$$v_{i,n}^{-1}[A] = (\zeta_{i,n} \circ v_i)^{-1}[A] = v_i^{-1}[\zeta_{i,n}^{-1}[A]] = v_i^{-1}[E].$$

With this we can conclude that $v_i^{-1}[E] \in \mathcal{B}_{\mathcal{X},i}^{ob}$. □

Because the identity morphism is the unique morphism from the universal choice structure to itself it follows as a simple corollary from Proposition 6 that $\mathcal{B}_{\mathcal{U},i}^{ob}$ is equal to the algebra \mathcal{B}_i of Ω_i . This property has already been observed in other settings and is called ‘‘minimality’’ in [DT08].

Before we provide the proof of Proposition 1 we would like to mention that the proposition shows that non-redundancy is the same as the notion of observability from the theory of coalgebras. In [Jac16] a coalgebra is called ‘‘observable’’ if the unique morphism from the coalgebra to the terminal coalgebra is monic. In [Rut00] such coalgebras are called ‘‘simple’’.

Proof of Proposition 1. For the direction from left to right assume that \mathcal{X} is non-redundant and consider the unique morphism $v = (v_i, v_j)$ from $\mathcal{X} = (T_i, T_j, \theta_i, \theta_j)$ into the universal choice structure \mathcal{U} . To show that v_i is injective consider $t, t' \in T_i$ such that $t \neq t'$. Because \mathcal{X} is non-redundant $\mathcal{B}_{\mathcal{X},i}^{ob}$ is separating, and thus there is some $E \in \mathcal{B}_{\mathcal{X},i}^{ob}$ with $t \in E$ and $t' \notin E$. By Proposition 6 there is then some E' that is measurable in Ω_i such that $E = v_i^{-1}[E']$. But then it must be the case that $v_i(t) \neq v_i(t')$, because otherwise $v_i(t') = v_i(t) \in E'$ and thus $t' \in v_i^{-1}[E'] = E$. The same reasoning works for j in place of i .

For the other direction assume that the $v_i : T_i \rightarrow \Omega_i$ is injective. To show that $\mathcal{B}_{\mathcal{X},i}^{ob}$ is separating on T_i consider any $t, t' \in T_i$ with $t \neq t'$. Because v_i is injective we have $v_i(t) \neq v_i(t')$. The algebra on Ω_i , which is defined as the limit algebra of the cochain $(\Omega_{i,n}, \xi_{i,n})_{n \in \omega}$ for discrete spaces $\Omega_{i,n}$, is separating because any two distinct infinite sequences have distinct projections to some finite level n , where they can be separated. Thus there must be a measurable set E in Ω_i such that $t \in E$ and $t' \notin E$. By Proposition 6 the set $v_i^{-1}[E]$ is in $\mathcal{B}_{\mathcal{X},i}^{ob}$ and obviously this set separates t and t' . The same argument can be used with j in place of i . □

E Proofs for Section 4

E.1 Preliminary observations

It is easy to check that \mathcal{P} , and hence also Π and Π^\heartsuit are functors, that is, they preserve identities and composition. The construction of the universal preference structure can then be carried out analogously to the construction for the universal choice structure given above. In fact it is only required to reprove a variant of Theorem 5 for the functor \mathcal{P} , which is done implicitly in the proofs of Section 3 from [DT08].

E.2 Proposition 2

We now argue that the map $m_X : \mathcal{P}X \rightarrow \mathcal{C}X$ is injective at every set X . To show this, assume we have two preference relations \preceq and \preceq' over X such that $m_X(\preceq) = m_X(\preceq')$. We need to argue that $x \preceq x'$ iff $x \preceq' x'$ for all $x, x' \in X$. Since the situation is symmetric it suffices to check one direction. Hence assume that $x \preceq x'$. We want to show that $x \preceq' x'$. It suffices to consider the case where $x \neq x'$ because otherwise $x \preceq' x'$ follows because \preceq' is reflexive.

As $x \neq x'$ and $x \preceq x'$ it follows that it can not be the case that $x' \preceq x$, because otherwise there would be a contradiction with the anti-symmetry of \preceq . As we explain in Remark 1 this use of anti-symmetry is crucial. As $x \preceq x'$ and not $x' \preceq x$ it follows that x' is the only maximal element of the set $\{x, x'\}$ in the relation \preceq . Hence $m_X(\preceq)(\{x, x'\}) = \{x'\}$.

By the assumption that $m_X(\preceq) = m_X(\preceq')$ it follows that $m_X(\preceq')(\{x, x'\}) = \{x'\}$. But this is only possible if $x \preceq' x'$, which is what we had to show.

E.3 Proposition 3

Proposition 3 states that λ is a natural transformation from the functor Π to the functor Γ . It is easy to check that this reduces to the claim that m is a natural transformation from \mathcal{P} to \mathcal{C} . This means that we need to show that for every function $f : X \rightarrow Y$ it holds that $m_X \circ \mathcal{P}f = \mathcal{C}f \circ m_Y$. Note that here the Y and X are swapped because \mathcal{P} and \mathcal{C} are contravariant functors.

Fix any function $f : X \rightarrow Y$, preference relation $\preceq \in \mathcal{P}Y$ and finite set $K \subseteq X$. We have to show that

$$m_X(\mathcal{P}f(\preceq))(K) = \mathcal{C}f(m_Y(\preceq))(K).$$

For the left-to-right inclusion take any $x \in m_X(\mathcal{P}f(\preceq))(K)$. We need to show that $x \in \mathcal{C}f(m_Y(\preceq))(K)$. This means we want to show that $x \in f^{-1}[m_Y(\preceq)(f[K])] \cap K$. Our assumption that $x \in m_X(\mathcal{P}f(\preceq))(K)$ means that x is a maximal element of the set K in the order $\preceq^f = \mathcal{P}f(\preceq)$. Hence $x \in K$ and it remains to show that $x \in f^{-1}[m_Y(\preceq)(f[K])]$, which means that $f(x)$ is a maximal element of the set $f[K]$ in the order \preceq . So consider any other element of $f[K]$, which must be of the form $f(x')$ for some $x' \in K$, and assume that $f(x) \preceq f(x')$. We need to show that then also $f(x') \preceq f(x)$. From equation (1) in Section 4.1 defining \preceq^f it follows that $x \preceq^f x'$. Then we can use that x is a maximal element in \preceq^f to conclude that $x' \preceq^f x$. Using (1) again, we then obtain the required $f(x') \preceq f(x)$.

Now consider the right-to-left inclusion. Take any $x \in f^{-1}[m_Y(\preceq)(f[K])] \cap K$. We need to show that $x \in m_X(\mathcal{P}f(\preceq))(K)$, which means that x is a maximal element of the set K in the order $\preceq^f = \mathcal{P}f(\preceq)$. Clearly $x \in K$. To show that x is a \preceq^f -maximal element in K pick any other x' in K with $x \preceq x'$. We need to show that $x' \preceq x$. From $x \preceq x'$ it follows with (1) that $f(x) \preceq f(x')$. We now use that $x \in f^{-1}[m_Y(\preceq)(f[K])]$. From this it follows that $f(x) \in m_Y(\preceq)(f[K])$. This means that $f(x)$ is maximal in the set $f[K]$. Because $x' \in K$ it holds that also $f(x')$ is in the set $f[K]$. By the maximality of $f(x)$ in $f[K]$ it follows from $f(x) \preceq f(x')$ that $f(x') \preceq f(x)$. With (1) we obtain $x' \preceq x$.

E.4 Theorem 3

In the proof of Theorem 3 we need a further concept from category theory which is a generalization of monic morphism. A family of morphism $(f_j : Y \rightarrow X_j)_{j \in J}$ for any index set J is *jointly monic* if for all further morphisms $g, h : T \rightarrow Y$ it holds that if $f_j \circ g = f_j \circ h$ for all $j \in J$ then already $g = h$.

Families of monic morphism are closely related to limits of cochains. Using the universal property of the limit X_ω with projections ζ_n of a cochain $(X_n, f_n)_{n \in \omega}$, it is easy to show that the family of all projections $(p_n)_{n \in \omega}$ is jointly monic. Moreover, if we have another object T with a jointly monic family $(g_n : T \rightarrow X_n)_{n \in \omega}$ such that $g_n = f_n \circ g_{n+1}$ for all n then the unique morphism $u : T \rightarrow X_\omega$ that exists because of the universal property of the limit X_ω is monic.

To prove Theorem 3, let $\mathcal{U}' = (\Omega', \mu')$ be the universal preference structure, presented as coalgebra for Π^\heartsuit . We assume that $\Omega', \mu', \Omega'_n, \zeta'_n, \dots$ are defined in the same way as the objects $\Omega, \mu, \Omega_n, \zeta_n, \dots$ are defined in Section D for the universal choice structure, just using Π^\heartsuit instead of Γ^\heartsuit .

Then consider the choice structure $\lambda(\mathcal{U}') = (\Omega', \lambda_{\Omega'}^\heartsuit \circ \mu')$ where λ^\heartsuit is a natural transformation from Π^\heartsuit to Γ^\heartsuit that is defined such that it applies λ componentwise. Note that this definition of $\lambda(\mathcal{U}')$ corresponds to the one given in Section 4.3. Let v be the unique morphism from $\lambda(\mathcal{U}')$ to the universal choice structure \mathcal{U} that exists according to Theorem 2.

The claim of Theorem 3 is that this v is monic. Because in Theorem 2 v was obtain using the universal property of the limit Ω from the family of approximations $(v_n : \Omega'_n \rightarrow \Omega_n)_{n \in \omega}$ it suffices to show that this family is jointly monic.

Define morphisms $\delta_0 = !_{\Omega'_0} : \Omega'_0 \rightarrow \Omega_0$ and inductively $\delta_{n+1} = \Gamma^\heartsuit \delta_n \circ \lambda_{\Omega'_n}^\heartsuit : \Omega'_{n+1} \rightarrow \Omega_{n+1}$. One can show with an induction over n that all these δ_n are monic. The base case this holds because Ω'_0 is the terminal object of \mathbf{Unc}^2 and in the inductive step we use Proposition 2 and the fact that Γ^\heartsuit preserves monics, which we show in Section C.2. The latter needs that the image of the injective measurable function $\delta_n : \Omega'_n \rightarrow \Omega_n$ that is preserved has the discrete algebra. This is the case because one can show that if A_i and A_j are finite then so are all the Ω'_n .

We then prove by induction on n that

$$v_n = \delta_n \circ \zeta'_n. \quad (11)$$

It follows that the v_n are jointly monic because the ζ'_n are projections out of the limit Ω' and hence jointly monic and the δ_n are all monic.

For the base case, of (11), we have that $v_0 = \delta_0 \circ \zeta'_0$ because both morphism map to the terminal object Ω_0 of \mathbf{Unc}^2 .

For the inductive step we use the following computation:

$$\begin{aligned} v_{n+1} &= \Gamma^\heartsuit v_n \circ \lambda_{\Omega'_n}^\heartsuit \circ \mu' && \text{definition of } v_{n+1} \\ &= \Gamma^\heartsuit (\delta_n \circ \zeta'_n) \circ \lambda_{\Omega'_n}^\heartsuit \circ \mu' && \text{induction hypothesis} \\ &= \Gamma^\heartsuit \delta_n \circ \Gamma^\heartsuit \zeta'_n \circ \lambda_{\Omega'_n}^\heartsuit \circ \mu' && \Gamma^\heartsuit \text{ functor} \\ &= \Gamma^\heartsuit \delta_n \circ \lambda_{\Omega'_n}^\heartsuit \circ \Pi^\heartsuit \zeta'_n \circ \mu' && \lambda^\heartsuit \text{ natural transformation} \\ &= \Gamma^\heartsuit \delta_n \circ \lambda_{\Omega'_n}^\heartsuit \circ \tau'_{n+1} \circ \mu' && \text{definition of } \tau'_{n+1} \\ &= \Gamma^\heartsuit \delta_n \circ \lambda_{\Omega'_n}^\heartsuit \circ \zeta'_{n+1} && \text{uniqueness property of } \mu' \\ &= \delta_{n+1} \circ \zeta'_{n+1} && \text{definition of } \delta'_{n+1} \end{aligned}$$