Polynomial Conjoint Measurement

Introduction

In the previous chapter, we studied additive conjoint measurement.

We had some relational structure $\langle A_1 \times \ldots \times A_n, \succsim \rangle$ such that for each A_i we could find a $\phi_i : A_i \to \mathbb{R}$ for each $i = 1, \ldots, n$ such that for all $a_i, b_i \in A_i$,

$$a_1 \ldots a_n \succsim b_1 \ldots b_n \text{ iff } \sum_{i=1}^{n} \phi_i(a_i) \geq \sum_{i=1}^{n} \phi_i(a_i)$$

Introduction

In this chapter, we're interested in the more general case:

We have some relational structure $\langle A_1 \times \ldots \times A_n, \succsim \rangle$ and we want to find a $\phi_i : A_i \to \mathbb{R}$ for each A_i and a polynomial $F : \mathbb{R}^n \to \mathbb{R}$ such that for all $a_i, b_i \in A_i$,

$$a_1 \ldots a_n \succsim b_1 \ldots b_n \text{ iff}$$
$$F[\phi_1(a_1), \ldots, \phi_n(a_n)] \geq F[\phi_1(b_1), \ldots, \phi_n(b_n)]$$

This subsumes the additive case where $F(x_1, \ldots, x_n) = \sum_{i=1}^{n} x_n$.

Introduction

The results we'll obtain will allow us to classify some empirical structures corresponding to a small class of particularly well behaved polynomials.

Before we get to those results, we're going to talk about a property called decomposability, and its relationship to polynomial measurement.

Defining Decomposability

There's a natural condition we might want of our n-factor structures: that we might be able to obtain a real numbered representation of each of the n-factors, and then construct a function from $\mathbb{R}^n$ to $\mathbb{R}$ that preserves the ordering of the empirical structure in the reals.

This section is about *decomposability*, which is slight strengthening of that natural condition.

Defining Decomposability

Definition

A structure $\langle A_1 \times \ldots \times A_n, \succsim \rangle$ is called *decomposable* when there are functions $\phi_i : A_i \to \mathbb{R}$ for $i = 1, \ldots, n$ and a function $F : \mathbb{R}^n \to \mathbb{R}$ that's one-to-one in each variable separately, such that for all $a, b \in A$,

$$a \succsim b \text{ iff } F[\phi_1(a_1), \ldots, \phi_n(a_n)] \geq F[\phi_1(b_1), \ldots, \phi_n(b_n)]$$

Definition

A structure is *monotonically decomposable* when, in addition to being decomposable, the associated F is strictly increasing in each variable separately.

Defining Decomposability

One wording I'll use that the book doesn't is that if $\langle A_1 \times \ldots \times A_n, \succsim \rangle$ is a decomposable structure and F is a function that meets the criteria above, I'll say that F *decomposes* $\langle A_1 \times \ldots \times A_n, \succsim \rangle$.

F is One-to-one in Each Variable Separately

To get back to the definition, let's clarify what "one-to-one in each variable separately" means.

We say F is one-to-one in each variable separately when $F(y_1, \ldots, y_{i-1}, x_i y_{i+1}, \ldots, y_n) = F(y_1, \ldots, y_{i-1}, x_i' y_{i+1}, \ldots, y_n)$ implies $x_i = x_i'$ for all $i = 1, \ldots, n$.

If we keep everything the same, but change one input variable to a different value, the function's output will be different, no matter what we change it to.

This doesn't mean that F is necessarily one-to-one overall: we don't have any requirements about what happens when we change multiple variables at once.

The Constraints on F

So there are two restrictions placed on F:

- It has to be one-to-one in each variable separately
- It has to (together with the ϕ_i functions) preserve the empirical structure's order in the reals.

Decomposability and Polynomial Representation

As it turns out, the "one-to-one in each variable separately" condition is non-trivial, and represents the "strengthening" I mentioned a moment ago.

Predecomposability is what I'll call the property which is the same as decomposability, but without the condition that F must be one-to-one in each variable.

The set of structures for which there exists a polynomial model is just a subset of the predecomposable structures, but it turns out that there are structures for which there are polynomial models that aren't decomposable (and vice versa.)

Decomposability and Polynomial Representation

So decomposability is not a necessary or sufficient condition for having a polynomial representation. As it turns out, we'll generally have to tweak the domain of the polynomials we're interested in, in order to get them to satisfy decomposability.

What does decomposability bring to the table?

- It ensures that, disregarding values that are always equivalent, each of the system's factors always has an effect on the final product.

Necessary and Sufficient Conditions

Theorem

Theorem 1. $\langle A_1 \times \ldots \times A_n, \succsim \rangle$ *is decomposable iff:*

- $\succsim$ *is a weak order.*

$\langle A, \succsim \rangle$ is necessarily a weak order. The reals are a weak order, so the $a \succsim b$ iff $F[\phi_1(a_1), \ldots, \phi_n(a_n)] \geq F[\phi_1(b_1), \ldots, \phi_n(b_n)]$ condition means that any two a and b have to be comparable, and the transitivity has to hold.

Necessary and Sufficient Conditions

Theorem

Theorem 1. $\langle A_1 \times \ldots \times A_n, \succsim \rangle$ *is decomposable iff:*

- $\succsim$ *is a weak order.*
- $A/\sim$ *has a countable order-dense subset.*

$A/\sim$, the set of equivalence classes of A under $\sim$, has a countable order-dense subset. This comes from Theorem 2.2, which says that a simple order has a countable order-dense subset iff there's an injective homomorphism from it to the reals.

Necessary and Sufficient Conditions

Theorem

Theorem 1. $\langle A_1 \times \ldots \times A_n, \succsim \rangle$ *is decomposable iff:*

- $\succsim$ *is a weak order.*
- $A/\sim$ *has a countable order-dense subset.*
- $\sim$ *satisfies substitutability.*

Substitutability comes from the one-to-one in each variable condition. We say $\langle A_1 \times \ldots \times A_n, \succsim \rangle$ satisfies substitutability iff for any choice of the involved variables, $b_1 \cdots b_{i-1} \mathbf{a_i} b_{i+1} \cdots b_n \sim b_1 \cdots b_{i-1} \mathbf{a'_i} b_{i+1} \cdots b_n$ iff $c_1 \cdots c_{i-1} \mathbf{a_i} c_{i+1} \cdots c_n \sim c_1 \cdots c_{i-1} \mathbf{a'_i} c_{i+1} \cdots c_n$. So we can hold a_i constant and change variables on both sides of the $\sim$, and the relation will still hold.

Independence and Monotonic Decomposability

There's another condition worth defining:

If F is strictly increasing in each component then for any choice of the involved variables, $b_1 \cdots b_{i-1}\mathbf{a_i}b_{i+1} \cdots b_n \succsim b_1 \cdots b_{i-1}\mathbf{a_i'}b_{i+1} \cdots b_n$ iff $c_1 \cdots c_{i-1}a_i c_{i+1} \cdots c_n \succsim c_1 \cdots c_{i-1}\mathbf{a_i'}c_{i+1} \cdots c_n$. Here, the order $\succsim_i$ induced on A_i by fixing all the non-A_i components is independent of what values we choose to fix those components at.

In this case, we say that A_i is *independent* of $\times_{j \neq i} A_j$.

Theorem

A structure $\langle A_1 \times \ldots \times A_n, \succsim \rangle$ is monotonically decomposable iff it's decomposable and each A_i is independent of $\times_{j \neq i} A_j$.

I Can't Believe It's Not a Representation Theorem

What we have so far is almost like a normal representation theorem, but not quite.

What we're asserting here is that iff the above conditions hold, we can map $A_1 \times \ldots \times A_n$ to $\mathbb{R}$, by way of some intermediate functions $\phi_i : A_i \to \mathbb{R}$ and a one-to-one in each variable separately function $F : \mathbb{R}^n \to \mathbb{R}$, such that $\succsim$ is preserved as $\geq$ under ϕ.

Normally, we specify the function F in the representation, and most of the time, it's just addition. Here, we leave the F unspecified, and much of the rest of the chapter consists of investigating a few suitable F functions.

Uniqueness

There are two uniqueness questions worth bringing up:

1. How much can we fiddle around with our ϕ_i functions (while retaining the same F) such that we've still got a homomorphism?
2. How much can we fiddle around with both our ϕ_i functions and F, such that the homomorphism is preserved?

The answer to the former question depends on the F, so we'll come back to that. We can deal with the latter question right now.

Uniqueness

Suppose we have a decomposing homomorphism $\phi : A_1 \times \ldots \times A_n \to \mathbb{R}$ made up of ϕ_i functions and a suitable F, and we also have some real valued, strictly increasing function $h : \text{ran}(F) \to \mathbb{R}$ (the domain could probably be shrunk).

Strictly increasing functions preserve order, so $\phi' = h \circ \phi$ is a decomposing homomorphism as well, and we can find corresponding, modified versions of F and ϕ_i (call them F' and ϕ_i') to go along with it.

Conversely, if F and F' both decompose $\langle A_1 \times \ldots \times A_n, \succsim \rangle$ then there's a strictly increasing h, that takes one to the other, and the corresponding ϕ_i functions will be constrained by that h (but not generally uniquely determined).

Nice Polynomials

Polynomials aren't necessarily strictly increasing or even one-to-one in each variable. Therefore, not every polynomial can decompose some structure.

But many polynomials almost fit the bill, and we can tweak them so that they do.

Suppose we have some polynomial F. If all the coefficients in F are positive and the domain of each varible in F is strictly positive, then F is strictly increasing in each variable.

I'll call a polynomial that satisfies these conditions *nice*.

If a structure satisfies a nice polynomial model, it's monotonically decomposable.

Equivalence for Nice Polynomials

Suppose a structure satisfies two different nice polynomial models $F(\phi_1(a_1), \ldots, \phi_n(a_n))$ and $F'(\phi'_1(a_1), \ldots, \phi'_n(a_n))$.

In this case there are strictly increasing functions h and h_i $(i = 1, \ldots, n)$ defined on the positive reals such that $F'(x_1, \ldots, x_n) = h(F(h_1^{-1}(x_1), \ldots h_n^{-1}(x_n)))$ and we say that the two polynomial models are *equivalent*.

Equivalence for Nice Polynomials

One question we can ask is: given some nice polynomial, when can we find such h and h_i functions that will take it into an equivalent nice polynomial?

We don't have a complete answer. If h and h_i^{-1} are nice polynomials, that's sufficient, but we can find some examples where non-polynomial h and h_i^{-1} functions will do the trick.

Polynomials and Decomposability

If our polynomials aren't so nice, for instance if they're defined for non-positive values of the variables, they're not necessarily going to satisfy the decomposability conditions right out of the box.

The example given in the book is $x_1 x_2$ defined on $\mathbb{R}^2$. It's not one-to-one in each variable, (consider the case where x_1 is set to 0), so it can't be a representation of a decomposable structure.

However, suppose we exclude zeroes from its domain, i.e. we define it over $(\mathbb{R} - \{0\}) \times (\mathbb{R} - \{0\})$. In that case, the functions becomes one-to-one in each variable, and it could satisfy some decomposable structure.

Polynomials and Decomposability

Also, notice that $x_1 x_2$ isn't strictly increasing in each variable if we let the variables take on negative values.

However, it is either strictly increasing or strictly decreasing in each variable, depending on the sign of the other variable, which means it satisfies a condition akin to monotone decomposability. We'll investigate these properties more in a bit.

We can also define equivalent polynomial models in the case where nonpositive values are permitted, that turns out more complicated. (The domain ends up giving us problems in a lot of cases, and we can find weird situations like where two polynomials are equivalent for any finite structure but not for infinite ones.)

Defining Simple Polynomials

We're going to define a class of polynomials which we'll call *simple*. A polynomial is simple if it can be recursively split into either products or sums of smaller polynomials with no variables in common.

Here's the recursive definition of $S(X)$, the simple polynomials in $X = \{x_1, \ldots, x_n\}$.

Definition

$S(X)$ is the smallest set of polynomials such that:

- $x_i \in S(X)$ for $i = 1, \ldots, n$
- If Y_1 are disjoint, non-empty subsets of X and $F_1 \in S(Y_1)$, $F_2 \in S(Y_2)$ then $F_1 + F_2$ and $F_1 F_2$ are in $S(X)$.

Simple Polynomials in Three Variables

There are effectively only four kinds of simple polynomials with exactly three variables:

Additive: $x_1 + x_2 + x_3$

Distributive: $(x_1 + x_2)x_3$, $(x_3 + x_2)x_1$, $(x_1 + x_3)x_2$

Dual-distributive: $x_1x_2 + x_3$, $x_1x_3 + x_2$, $x_2x_3 + x_1$

Multiplicative: $x_1x_2x_3$

Strictly speaking, there are four more polynomials that can be formed by permuting the positions of the variables in the distributive or dual-distributive polynomials, but since we can order our factors any way we want, this doesn't really matter.

Simple Polynomials in Four Variables

The four variable simple polynomials can be formed from the three variable ones to get ten different kinds:

Four kinds by simply tacking on another addition term to a three variable form, four kinds by multiplying one of the three variable forms by a new variable and two more polynomials in new forms: $(x_1 + x_2)(x_3 + x_4)$ and $x_1x_2 + x_3x_4$.

For the rest of this chapter, we'll be concerned exclusively with simple polynomials.

Notation

Definitions

Suppose $N \subseteq \{1, \ldots, n\}$.

- $a^{(N)}$ is a vector with components taken from each $A_{i \in N}$.
- $A^{(N)}$ is the set of all such $a^{(N)}$ vectors.

Generalizing Independence

Suppose we have some structure $\langle A_1 \times \ldots \times A_n, \succsim \rangle$ and also some sets N_1, N_2 and N_3 which together form a partition of the index set $\{1, \ldots, N\}$.

Definition

- I'll denote the element of $A_1 \times \ldots \times A_n$ composed of elements from the vectors $a^{(N_1)}$, $a^{(N_2)}$, $a^{(N_3)}$ as $a^{(N_1)}a^{(N_2)}a^{(N_3)}$.

Generalizing Independence

Suppose we have some structure $\langle A_1 \times \ldots \times A_n, \succsim \rangle$ and also some sets N_1, N_2 and N_3 which together form a partition of the index set $\{1, \ldots, N\}$.

We can say $A^{(N_1)}$ is independent of $A^{(N_2)}$ if given some $a^{(N_3)}$:

$a^{(N_1)}a^{(N_2)}a^{(N_3)} \succsim b^{(N_1)}a^{(N_2)}a^{(N_3)}$ iff
$a^{(N_1)}c^{(N_2)}a^{(N_3)} \succsim b^{(N_1)}c^{(N_2)}a^{(N_3)}$ for any choice of $a^{(N_1)}$, $b^{(N_1)}$ and $c^{(N_2)}$.

Sign Dependence

We say $A^{(N_1)}$ is sign dependent on $A^{(N_2)}$ if $A^{(N_2)}$ can be partitioned into three sets, $S^+(N_1, N_2)$, $S^0(N_1, N_2)$ and $S^-(N_1, N_2)$, such that:

- $A^{(N_1)}$ is independent of each
- $S^+(N_1, N_2) \cup S^-(N_1, N_2)$ is non-empty
- The relation induced on $A^{(N_1)}$ by elements from $S^+(N_1, N_2)$ is the converse of that induced by elements from $S^-(N_1, N_2)$
- The relation induced on $A^{(N_1)}$ by elements from $S^0(N_1, N_2)$ is degenerate.

If $S^0(N_1, N_2)$ and exactly one of $S^+(N_1, N_2)$ and $S^-(N_1, N_2)$ are empty, then $A^{(N_1)}$ is independent from $A^{(N_2)}$. If two or three of the sets are non-empty, then $A^{(N_1)}$ is properly sign dependent on $A^{(N_2)}$.

Outline of the rest of the chapter

We use sign dependence to describe which factors of the simple polynomials are dependent on one another.

We can use this information to narrow down what simple polynomials are compatible with the empirical structure.

If all the factors are independent, it could technically be any of the simple polynomials, so we use some joint independence conditions.

There's also a distributive cancellation condition which is necessary for a distributive representation.

Eventually we're presented with a flow chart for diagnosing the proper simple polynomial.

Finally, we end up with representation and uniqueness theorems for multiplicative, distributive and dual-distributive polynomials.

Multiplicative Case

Theorem

Suppose that $\succsim$ *is a binary relation on* $A = A_1 \times A_2 \times A_3$ *for which the following axioms are satisfied:*

- $\succsim$ *is a weak order.*
- *Each pair of factors is sign dependent on the third.*
- *Every strictly bounded standard sequence in one factor is finite.*
- *Unrestricted solvability holds.*

Then, there exist real-valued functions ϕ_i *on* A_i, $i = 1, 2, 3$ *such that, for all* $a, b \in A$,
$a \succsim b$ *iff* $\phi_1(a_1)\phi_2(a_2)\phi_3(a_3) \geq \phi_1(b_1)\phi_2(b_2)\phi_3(b_3)$
Moreover, real-valued functions satisfying this property are unique up to the transformations:

$$\phi_i(a_i) \rightarrow \begin{cases} \alpha_i(\phi_i(a_i))^\beta & \text{if } \phi_i(a_i) \geq 0 \\ -\alpha_i(-\phi_i(a_i))^\beta & \text{if } \phi_i(a_i) \leq 0 \end{cases}$$

where α_i *and* β *are real numbers such that* $\beta > 0$ *and* $\alpha_1\alpha_2\alpha_3 > 0$.

Distributive Case

Theorem

Suppose that $\succsim$ *is a binary relation on* $A = A_1 \times A_2 \times A_3$ *for which the following axioms are satisfied:*

- $\succsim$ *is a weak order.*
- $A_1 \times A_2$ *and* A_3 *are mutually sign dependent.*
- $\langle A_1 \times A_2, A_3, \sim\rangle$ *satisfies the Thomsen condition of Definition 4.*
- *Distributive cancellation holds.*
- *For any induced ordering on* $A_1 \times A_2$*, every strictly bounded standard sequence in one factor is finite.*
- *Unrestricted solvability holds.*
- $(A_1 \times A_2)^0$ *and* $(A_3)^0$ *are nonempty.*

Then, there exist real-valued functions ϕ_i *on* A_i*,* $i = 1, 2, 3$ *such that, for all* $a, b \in A$*,*
$a \succsim b$ *iff* $(\phi_1(a_1) + \phi_2(a_2))\phi_3(a_3) \geq (\phi_1(b_1) + \phi_2(b_2))\phi_3(b_3)$
Moreover, real-valued functions satisfying this property are unique up to the transformations:

- $\phi_1 \to \alpha\phi_1 + \beta$
- $\phi_2 \to \alpha\phi_2 - \beta$
- $\phi_3 \to \gamma\phi_3$

where α, β *and* γ *are real numbers such that* $\alpha\gamma > 0$.

Dual-Distributive Case

Theorem

Suppose that $\succsim$ is a binary relation on $A = A_1 \times A_2 \times A_3$ for which the following axioms are satisfied:

- *$\succsim$ is a weak order.*
- *$A_1 \times A_2$ and A_3 are mutually independent, while A_2 and A_1 are mutually sign dependent.*
- *$\langle A_1 \times A_2, A_3, \sim \rangle$ satisfies the Thomsen condition of Definition 4.*
- *Dual-distributive cancellation holds.*
- *Regarding $\langle A_1 \times A_2, A_3, \sim \rangle$ as a two-component structure, each component has the property that every strictly bounded standard sequence is finite.*
- *Unrestricted solvability holds.*
- *$(A_1)^0$ and $(A_2)^0$ are nonempty.*

Then, there exist real-valued functions ϕ_i on A_i, $i = 1, 2, 3$ such that, for all $a, b \in A$,
$a \succsim b$ iff $\phi_1(a_1)\phi_2(a_2) + \phi_3(a_3) \geq \phi_1(b_1)\phi_2(b_2) + \phi_3(b_3)$
Moreover, real-valued functions satisfying this property are unique up to the transformations:

- $\phi_1 \to \alpha_1 \phi_1$
- $\phi_2 \to \alpha_2 \phi_2$
- $\phi_3 \to (\alpha_1 \alpha_2)\phi_3 + \beta$

where α_1, α_2 and β are real numbers such that $\alpha_1 \alpha_2 > 0$.

Conditional Expected Utility

Conditional Expected Utility

We're concerned with modeling subjective perceptions of utility associated with decisions, where the subject has some control over the possible outcomes of a chance set up and the consequences associated with those outcomes.

Conditional Expected Utility

There are three stages to the situation being modeled:

Decision We have some subject who has a number of options available to her. The subject makes a choice which constrains the possible subsequent things that can occur in the chance set up.

Conditional Expected Utility

There are three stages to the situation being modeled:

Decision We have some subject who has a number of options available to her. The subject makes a choice which constrains the possible subsequent things that can occur in the chance set up.

Outcome A particular thing occurs, usually determined by chance (or at least by a process which appears to involve chance to the subject.)

Conditional Expected Utility

There are three stages to the situation being modeled:

Decision We have some subject who has a number of options available to her. The subject makes a choice which constrains the possible subsequent things that can occur in the chance set up.

Outcome A particular thing occurs, usually determined by chance (or at least by a process which appears to involve chance to the subject.)

Consequence There is some consequence for the subject associated with the above outcome, determined by the choice that the subject made.

Representation of Consequences

All the consequences we're interested in for the consequence stage are drawn from a set C. The members of C can be arbitary consequences like getting a book, losing five dollars, feeling happy or summary execution.

Representation of Outcomes

Next we're concerned with what can happen in the middle stage. We start with a set X of outcomes that could happen as a result of whatever chance process we're interested in.

We can use X to define a few other sets we're interested in. Let E be a nonempty set of subsets of X that's closed under complement and finite union (i.e., it's an algebra over X.) We'll call the sets in E events, and we'll be associating probabilities with them later.

N is a subset of E: it's the set of "null" events, which shouldn't occur. When we're in the business of assigning probabilities to events, the events in N will have probability 0.

Representation of Decisions

Finally, we have a set of decisions called D.

The purpose of a decision is to determine both what possible outcomes can happen in the chance setup and how consequences are associated with those events, so we can represent a choice with a function $f_A : A \to C$, where $A \in E - N$.

The Empirical Relation

Now we can arrive at the empirical relation we're interested in:

A relation $\precsim$ can be derived over the decision set D by presenting subjects with pairs of decisions and determining which one they prefer.

The Empirical Relational Structure

All together, our empirical relational structure is made up of the following elements:

- X Elementary outcomes
- E Events (sets of outcomes)
- N Null events (impossible events)
- C Consequences
- D Decisions
- $\succsim$ A preference relation over the decisions

Things the Theory Doesn't Include

- An objective probability distribution on E

Things the Theory Doesn't Include

- An objective probability distribution on E
- Precisely what information the subject has about the decisions available to her

Why do it this way?

Question: Why do we need to model the intermediary stage, where we learn the result of the chance process?

We could just have a setup where the subject chooses between sets of possible consequences, with a conditional probability distribution over the consequences.

As it turns out, divorcing the chance outcomes and the consequences gives us a lot of representational power we wouldn't have otherwise.

For one thing, we'd have a hard time representing different decisions with the same possible consequences without formally including a probability model of some sort. Second, we'd like to have the possibility of not having utility strictly coupled to consequences, but rather having it as a function of all the factors involved in a decision.

This is really interesting, we'll get back to it later.

The Utility Function

We have our empirical relational structure $\langle X, E, N, C, D, \succsim \rangle$, we'd like to be able to represent it in the reals. The homomorphism from the decision set to the reals is going to be a utility function u.

$u(f_A)$ represents a numerical measure of utility assigned to the decision f_A. Notice, again, we associate utility with decisions most directly. In certain circumstances, we'll be able to treat $u(f_A)$ as the expected utility associated with the various consequences of f_A, but not in the most general case.

The Subjective Probability Measure

We expect that the utility the subject associates with a particular choice in which the possible outcomes come from the union of two (disjoint) events should be consistently weighted by her subjective impression of the relative probability of those events.

We'll represent this subjective probability formally in our representation by a function P which is defined over the event set E.

The Subjective Probability Measure

Since the decisions in D are just functions, if their domains are disjoint, we can take the set theoretic union of the functions, which gives us a well-defined function from the union of their domains to the union of their ranges.

So the union of f_A and g_B is $f_A \cup g_B(x) = \begin{cases} f_A(x) & \text{if } x \in A \\ g_B(x) & \text{if } x \in B \end{cases}$.

Notice that for $f_A \cup g_B$, we can end up with any of the consequences we got from f_A or g_B.

It would be useful to have a sense of which of A or B is more likely given that the outcome will be in $A \cup B$. As it turns out, the only probabilities we're really interested in are of the form $P(A|A \cup B)$ where A, B and $A \cup B$ are events which form the domain of some decisions f_A, g_B and $f_A \cup g_B$.

The Desired Representation

All together, the representation system we want involves an order preserving map $u : D \to \mathbb{R}$ and a subjective probability measure P on E such that:

- For all $R \in E$, $R \in N$ iff $P(R) = 0$
- For events $A, B \in E - N$, and $f_A, f_B \in D$:
  - $f_A \succsim g_B$ iff $u(f_A) \geq u(g_B)$
  - If A and B are disjoint then $u(f_A \cup g_B) = u(f_A)P(A|A \cup B) + u(g_B)P(B|A \cup B)$

So N represents subjectively impossible events, the utility function preserves subjective preference between decisions and the subjective utility of a union of disjoint decisions is weighted by the perceived conditional probability of the occurrence of outcomes from their domains.

The Axiom System

Closure Guarantees that D is sufficiently rich

- If $A, B \in E$ and $f_A, g_B \in D$, then:
  - If A and B are disjoint, then $f_A \cup g_B \in D$.
  - If $B \subset A$ then the restriction of f_A to B is in D.

The Axiom System

Closure Guarantees that D is sufficiently rich

Weak Order $\succsim$ is a weak ordering of D

The Axiom System

Closure Guarantees that D is sufficiently rich

Weak Order $\succsim$ is a weak ordering of D

Union Indifference A mix of two equivalent decisions is equivalent to both

- For disjoint $A, B \in E$ and $f_A, g_B \in D$, $f_A \sim g_B$ implies $f_A \cup g_B \sim f_A$.

The Axiom System

Closure Guarantees that D is sufficiently rich

Weak Order $\succsim$ is a weak ordering of D

Union Indifference A mix of two equivalent decisions is equivalent to both

Independence Adding/removing disjoint (sub)decisions doesn't affect preference order

- For disjoint $A, B \in E$ and $f_A^{(1)}, f_A^{(2)}, g_B \in D$, $f_A^{(1)} \succsim f_A^{(2)}$ iff $f_A^{(1)} \cup g_B \succsim f_A^{(2)} \cup g_B$.

The Axiom System

Closure Guarantees that D is sufficiently rich

Weak Order $\succsim$ is a weak ordering of D

Union Indifference A mix of two equivalent decisions is equivalent to both

Independence Adding/removing disjoint (sub)decisions doesn't affect preference order

Compatibility Two different utility interval orderings coincide

For disjoint $A, B \in E$, $f_A^{(i)} \sim g_B^{(i)}$ for $i = 1, 2, 3, 4$, $f_A^{(1)} \cup k_B^{(1)} \sim f_A^{(2)} \cup k_B^{(2)}$ and $h_A^{(1)} \cup g_B^{(1)} \sim h_A^{(2)} \cup g_B^{(2)}$, then $f_A^{(3)} \cup k_B^{(1)} \sim f_A^{(4)} \cup k_B^{(2)}$ iff $h_A^{(1)} \cup g_B^{(3)} \sim h_A^{(2)} \cup g_B^{(4)}$.

The Axiom System

Closure Guarantees that D is sufficiently rich

Weak Order $\succsim$ is a weak ordering of D

Union Indifference A mix of two equivalent decisions is equivalent to both

Independence Adding/removing disjoint (sub)decisions doesn't affect preference order

Compatibility Two different utility interval orderings coincide

Archimedian Condition ....

For disjoint $A, B \in E$, where N is a sequence of consecutive integers, $g_B^{(0)} \not\sim g_B^{(1)}$, $f_A^{(i)} \cup g_B^{(1)} \sim f_A^{(i+1)} \cup g_B^{(0)}$ for all, $i, i+1 \in N$ then either N is finite of $\{f_A^{(i)} | i \in N\}$ is unbounded.

The Axiom System

Closure Guarantees that D is sufficiently rich

Weak Order $\succsim$ is a weak ordering of D

Union Indifference A mix of two equivalent decisions is equivalent to both

Independence Adding/removing disjoint (sub)decisions doesn't affect preference order

Compatibility Two different utility interval orderings coincide

Archimedian Condition ....

Nullity Null events behave sanely

- If $R \in N$ and $S \subset R$, then $S \in N$.
- $R \in N$ iff, for all $f_{A \cup R} \in D$ where $R \in E$ and $A \in E - N$ are disjoint, $f_{A \cup R} \sim f_A$, where f_A is the restriction of $f_{A \cup R}$ to A.

The Axiom System

Closure Guarantees that D is sufficiently rich

Weak Order $\succsim$ is a weak ordering of D

Union Indifference A mix of two equivalent decisions is equivalent to both

Independence Adding/removing disjoint (sub)decisions doesn't affect preference order

Compatibility Two different utility interval orderings coincide

Archimedian Condition ....

Nullity Null events behave sanely

Nontriviality Nonnecessay, guarantees cancellation conditions

- $E - N$ has at least three pairwise disjoint elements and $D/\sim$ has at least two distinct equivalence classes.

The Axiom System

Closure Guarantees that D is sufficiently rich

Weak Order $\succsim$ is a weak ordering of D

Union Indifference A mix of two equivalent decisions is equivalent to both

Independence Adding/removing disjoint (sub)decisions doesn't affect preference order

Compatibility Two different utility interval orderings coincide

Archimedian Condition ....

Nullity Null events behave sanely

Nontriviality Nonnecessay, guarantees cancellation conditions

Restricted Solvability Nonnecessary solvability requirements

- Given $A, B \in E$ and $g_B \in D$, there's an $h_A \in D$ such that $h_A \sim g_B$.
- Given disjoint $A, B \in E$ and $h_A^{(1)} \cup g_B \succsim f_{A \cup B} \succsim h_A^{(2)} \cup g_B$, there's an $h_A \in D$ such that $h_A \cup g_B \sim f_{A \cup B}$.

The Representation Theorem

Theorem

If $\langle X, E, N, C, D, \succsim \rangle$ is a conditional decision structure, there exist real valued functions u on D and P on E such that $\langle X, E, P \rangle$ is a finitely additive probability space (see chapter 5) and for all $A, B \in E - N$, $R \in E$, $f_A, g_B \in D$,

- $R \in N$ *iff* $P(R) = 0$
- $f_A \succsim g_B$ *iff* $u(f_A) \geq u(g_B)$
- *If A and B are disjoint then* $u(f_A \cup g_B) = u(f_A)P(A|A \cup B) + u(g_B)P(B|A \cup B)$

Furthermore, P is totally unique and u is unique up to a positive linear transformation.

Some Definitions

Definition

If $c \in C$ and $A \in E - N$ then the function c_A such that $c_A(x) = c$ for any $x \in A$, is called a *constant decision function*.

Definition

A conditional decision $f_A \in D$ is a *gamble* if the image of f_A is finite and if for every c in the image of f_A, the set of elements mapped into c (i.e. $f_A^{-1}(c) = \{x | x \in A, f_A(x) = c\}$) is an event in $E - N$.

Gambles are the finite union of constant decisions, and the associated f_A^{-1} sets partition A.

Utility of Consequences

In the model that we've been developing so far we assign utilities to decisions, not directly to consequences.

We typically think of the consequences as being what is most directly valued by the subject.

We often want to assume the preference relation over decisions is determined exclusively by the values (and relative probabilities) the subject associates with the possible consequences.

Necessary Conditions

In order to be able to assign utilities directly to consequences, our structure has to satisfy two conditions:

- We have to be able to find a constant decision for each consequence.
- The constant decisions for each particular consequence must be preferentially equivalent.

In the most general case, neither of these conditions is guaranteed to be satisfied.

Utility of Consequences

Theorem

If $\langle X, E, N, C, D, \succsim \rangle$ is a conditional decision structure, such that for every $c \in C$:

1. *There's some $A \in E - N$ such that $c_A \in D$.*
2. *If $A, B \in E - N$ and $c_A, c_B \in D$ then $c_A \sim c_B$.*

By theorem 8.1, there exist utility and probability functions u and P associated with $\langle X, E, N, C, D, \succsim \rangle$.

Continued...

Utility of Consequences

Theorem, continued

Then there exists a well defined value function v *that gives us the utility associated with a consequence, such that for each* $c \in C$ *and* $c_A \in D$, $v(c) = u(c_A)$.

Every gamble $f_A \in D$ *is of the form* $f_A = \bigcup_{i=1}^{n} c_{A_i}^{(i)}$, *so we can calculate the utility of each gamble in terms of the value of its consequences:* $u(f_A) = \sum_{i=1}^{n} v(c^{(i)})P(A_i|A)$.

Other Possibilities

The property we've just been talking about is a nice one for a structure to have. I'll refer to structures that afford such a representation, as well as the representations themselves, as being *consequentially determined*.

When we're modeling conditional expected utility, we typically want the consequences that appear in our model to represent the situational outcomes that matter to the subject. Consequentially determined relations are ideal in that regard.

However, there are cases of interest that don't satisfy the restrictions above.

Relations Determined by Other Factors

For example, suppose we can find a constant decision for each consequence, but we don't require a consequence's constant decisions to be preferentially equivalent across different domain events.

One simple way we can model this is to have some function $w : (E - N) \to \mathbb{R}$ that associates a utility with the events that serve as the decision domains in a manner consistent with the subject's subjective probability assignments:

$$w(A \cup B) = w(A)P(A|A \cup B) + w(B)P(B|A \cup B)$$

In this case, if f_A is a gamble of the form $f_A = \bigcup_{i=1}^{n} c_{A_i}^{(i)}$, then $u(f_A) = w(A) + \sum_{i=1}^{n} v(c^{(i)})P(A_i|A)$.

Good Luck Getting This One Past IRB

Suppose we present two games of chance to a subject, and she's given the task of deciding which will occur. In both games, there are only two reasonably possible outcomes and the subject has no reason to suppose that they aren't equally likely in either case. As a consequence of the outcome of the games, the subject might either win a single nickel, or nothing at all, and again she supposes both are equally likely in either case.

Good Luck Getting This One Past IRB

Suppose we present two games of chance to a subject, and she's given the task of deciding which will occur. In both games, there are only two reasonably possible outcomes and the subject has no reason to suppose that they aren't equally likely in either case. As a consequence of the outcome of the games, the subject might either win a single nickel, or nothing at all, and again she supposes both are equally likely in either case.

The first game of chance is determined by a coin flip. (The subject wins a nickel if the coin lands on "heads".)

Good Luck Getting This One Past IRB

Suppose we present two games of chance to a subject, and she's given the task of deciding which will occur. In both games, there are only two reasonably possible outcomes and the subject has no reason to suppose that they aren't equally likely in either case. As a consequence of the outcome of the games, the subject might either win a single nickel, or nothing at all, and again she supposes both are equally likely in either case.

The first game of chance is determined by a coin flip. (The subject wins a nickel if the coin lands on "heads".)

The second game of chance is determined by the outcome of a game of Russian roulette between two of the subject's acquaintances.

Consequences

Here, the "consequences" as represented in the formal model have nothing to do with the things that determine the subject's decision preferences.

As noted above, when we're modeling conditional expected utility, we typically want the consequences in our model to represent the outcomes that matter to the subject.

The problem is that, there's no a priori way to determine what aspects of a situation are going to turn out to be relevant to a subject's preferences, so we can't guarantee that the way we choose to model the situation will allow for a consequence-determined preference relation.

Consequences

We can define a *perceived subjective consequence* as any outcome involved in the decision situation that the subject finds relevant in determining decision preferences.

In principle, supposing the subject's preferences are rational enough to fit into a conditional decision structure, given any decision situation, there's nothing that would seem to rule out the existence of a model in which the consequences in C describe the subject's perceived subjective consequences at least well enough to allow for an empirical preference relation such that admits a consequence-determined representation.

Thoughts about this?

Model Error

One possibility I find interesting involves using the representation above, with the additional $w(A)$ term in the utility function to, in some sense, measure model error.

The $w(A)$ term gives us an idea of the degree to which the subject's preferences are dependent on aspects of the situation associated with the decision domain events, rather than the associated consequences.

It doesn't tell us how to fix our model, but it gives us some ideas about where to look.

Expected Utility and Risk

In this situation, our D is a set of gambles with monetary consequences.

Definitions

$E(f_A)$ is conditional expectation associated with f_A
$V(f_A)$ is the variance associated with f_A

Expected Utility and Risk

Definitions

Preference order is VE-dependent whenever for all $f_A, g_B \in D$, $V(f_A) = V(g_B)$ and $E(f_A) = E(g_B)$ imply $f_A \sim g_B$.
Let $R(f_A) = \theta V(f_A) - (1 - \theta)E(f_A)$ where $0 < \theta \leq 1/$
Preference order is R-dependent whenever for all $f_A, g_B \in D$, $R(f_A) = R(g_B)$ implies $f_A \sim g_B$.
Any preference that's R-dependent is VE-dependent, but not conversely.

Theorem

Let $\langle X, E, N, C, D, \succsim \rangle$ be a conditional decision structure, with $C = \mathbb{R}$ and suppose that:

- *For all $n \in I^+$, $c^{(1)}, \ldots, c^{(n)} \in \mathbb{R}$, $p_1, \ldots, p_n \in \mathbb{R}^+$ (where $\sum_{i=1}^{n} p_i = 1$), there exist pairwise disjoint events $A_1, \ldots, A_n \in E - N$ with $P(A_i | \cup_{i=1}^{n} A_i) = p_i$ such that $c_{A_i}^{(i)} \in D$ for $i = 1, \ldots, n$. (We can divide the conditional probability between n constant functions any way we please.)*
- *For any $c \in C$, $A, B \in E - N$, if $c_A, c_B \in D$ then $c_A \sim c_B$ (constant decisions agree)*
- *$v(c) = u(c_A) = \sum_{m=0}^{\infty} \alpha_m c^m$ for all $c \in \mathbb{R}$ (the utility associated with a consequence is monotonically increasing)*

Then, a preference order is VE-dependent iff $v(c) = \alpha_0 + \alpha_1 c + \alpha_2 c^2$ and furthermore, there's no preference order that's R-dependent.

The Relation Between Subjective and Objective Probability

There's a brief section relating subjective and objective probability measures.

Earlier we put aside the objective probability distribution that determines the likelihood of each particular outcome's occurranceoccurrence, given a particular decision.

Now we'll briefly investigate the relation between our normal subjective probability measure P and the objective probability measure Q.

The Relation Between Subjective and Objective Probability

Theorem

Let P and Q be finitely additive probability measures on an algebra of sets E that are strictly increasing functions of each other. If for every pair of rational numbers r and s (where $r + s \leq 1$) there's some disjoint pair $R, S \in E$ such that $P(R) = r$ and $P(S) = s$, then $P = Q$.

If E is sufficiently rich, and $P \neq Q$, then P doesn't preserve the objective probability ordering of events given by Q.