Representing Suppositional Decision Theories with Sets of Desirable Gambles

University of Bristol

EPIMP

Kevin Blackwell

Motivating Example: Extortion

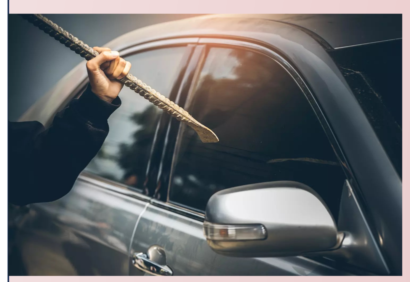

Suppose you have just parked in a seedy neighborhood when a man approaches and offers to "protect" your car from harm for $10. You recognize this as extortion and have heard that people who refuse "protection" invariably return to find their windshields smashed. Those who pay find their cars intact. You cannot park anywhere else because you are late for an important meeting. It costs $400 to replace a windshield. Should you buy "protection"? (example from James Joyce)

	Broken	Unbroken
Pay	-$410	-$10
Don't Pay	-$400	0

There seems to be a *dominance* argument against Paying: your windshield will either be Broken or Unbroken. In either case, Don't Pay provides strictly more value than Pay. … But this is silly.

Diagnosis: whether your windshield will be Broken or not *depends* on which act you perform.

Act-State Dependence

$\mathcal{X}$: the set of *states* of the world over which the agent doesn't have direct control; exactly one $X \in \mathcal{X}$ is true at the actual world.
$\mathcal{A}$: the set of *acts*; our agent will choose exactly one $A \in \mathcal{A}$.

Total outcome space: $\Omega \subseteq \mathcal{X} \times \mathcal{A}$. If acts and states are *logically* independent (any act-state pair is logically possible), $\Omega = \mathcal{X} \times \mathcal{A}$. We assume $\mathcal{X}$ and $\mathcal{A}$ are finite.

$u\colon \Omega \to \mathbb{R}$; we assume the agent has a utility function directly over total outcomes, indicating all preferences relevant to the decision. (In Extortion, the table above.)

Need a decision rule sensitive to the fact that the agent's beliefs about which state will obtain may be affected by their choice of act.

Imprecise SDTs and SDGs

Assume the agent has an imprecise prior $P \subseteq \mathbb{P}$.
$A \succ B$ iff $(\forall p \in P)\big(V_s(p,u,A) > V_s(p,u,B)\big)$.
Another obvious generalization (considered in the paper) would be to allow for imprecision about s, too.

A common decision rule in the act-state *independent* case is "supervaluation": $A \succ B$ iff $(\forall p \in P)\left(E_p(u_A) > E_p(u_B)\right)$, with $E_p(u_A) = \sum_{X \in \mathcal{X}} p(X)u(A,X)$; this decision rule for binary preferences is entailed by either Walley-Sen Maximality or E-Admissibility. Note here, the agent's credal set $P \subseteq \mathbb{P}(\mathcal{X})$ reflects only beliefs about states.

This can be nicely modeled with sets of desirable gambles. A gamble $g\colon \mathcal{X} \to \mathbb{R}$ represents a gain/loss to the agent determined by which *state* obtains; $\mathcal{L}(\mathcal{X})$ is the set of all such gambles; $D \subseteq \mathcal{L}$ is the agent's *set of desirable gambles* ($g \in D$ iff $g \succ 0$; g is preferred to the status quo).

1. Each act $A \in \mathcal{A}$ has a characteristic gamble $g_A(X) = u(A,X)$.
2. Link between D and P: $g \in D$ iff $(\forall p \in P)\big(E_p(g) > 0\big)$.
3. Read act preferences from gamble desirability: $A \succ B$ iff $g_A - g_B \in D$.
4. Act pricing: agent will pay ϵ for A iff $g_A - \epsilon \in D$; agent will sell A for ϵ iff $\epsilon - g_A \in D$.

In the *dependent* case, gambles must be defined on the total outcome space Ω; $g\colon \Omega \to \mathbb{R}$.

Suppositional Decision Theories

Well-studied in philosophy, but almost exclusively in a precise context.

Basic idea: evaluate an act A from the epistemic perspective of "supposing" A is the act you perform. Calculate the expected utility of each act from its perspective, choose whichever act maximizes suppositional expected utility.

Supposition rule: $s\colon \mathbb{P} \times \mathcal{P}_{\neg\emptyset}(\Omega) \to \mathbb{P}$. $s(p,R)$ maps a pair of probability function on Ω (interpretation: agent's precise prior) and nonempty event R (the supposed event) to a new probability function. One constraint: $s(p,R)(\omega) = 0$ for any $\omega \notin R$; supposition involves certainty in the supposed event. Otherwise, different s can represent a wide array of policies concerning how supposing R impacts other beliefs.

Choose the act which maximizes $V_s(p,u,\cdot) = E_{s(p,\cdot)}(u) = \sum_{\omega \in \Omega} s(p,\cdot)(\omega)u(\omega)$.

E.g.: Bayesian conditionalization yields Richard Jeffrey's Evidential Decision Theory (EDT). We can also represent various Causal Decision Theories in this formalism and some more exotic decision theories, like Functional Decision Theory.

Special Case: Generalized Imaging

Let $f\colon \mathcal{P}_{\neg\emptyset}(\Omega) \times \Omega \to \mathbb{P}$; we also require $f(R,\omega)(\omega') = 0$ for any $\omega' \notin R$. This is a "general imaging function." It's difficult to say much about their interpretation outside of particular decision theories.

We say a supposition rule is *representable* by generalized imaging when there's a general imaging function f such that $s(p,R)(\cdot) = \sum_{\omega \in \Omega} p(\omega)f(\omega,R)(\cdot)$, for every $p \in \mathbb{P}$ and every $R \in \mathcal{P}_{\neg\emptyset}(\Omega)$.

In the paper, we often only care whether supposition *of acts* is representable by generalized imaging.

Many versions of CDT supposition rules are representable by generalized imaging (including Pearl's do-operator); in this context, $f(R,\omega)(\omega')$ can be interpreted as a counterfactual probability: if the actual world is ω, the probability that ω' would (have) result(ed) from a causal intervention to make R obtain.

Can we use sets of desirable gambles to model supervaluated SDTs in a similar way?

We need to modify conditions 1 or 2, yielding two different kinds of representations (2 types of "bridges") linking an SDG to the judgments of the supervaluated SDT. Condition 4 (pricing) is also different between the 2 bridges.

Bridge Type 1

Maintain conditions 2-4, but use different characteristic gambles. **This kind of representation is possible iff supposition *of acts* is representable by generalized imaging.**

1. Each act $A \in \mathcal{A}$ has a characteristic gamble $g_A(\omega) = \sum_{\omega' \in \Omega} f(A,\omega')(\omega)u(\omega)$.
2. Link between D and P: $g \in D$ iff $(\forall p \in P)\big(E_p(g) > 0\big)$.
3. Read act preferences from gamble desirability: $A \succ B$ iff $g_A - g_B \in D$.
4. Act pricing: agent will pay ϵ for A iff $g_A - \epsilon \in D$; agent will sell A for ϵ iff $\epsilon - g_A \in D$.

Bridge Type 2

Use natural analogue of condition 1; maintain condition 3; use "effective" (fake) credal set for condition 2:
$P_{eff} = \left\{q \in \mathbb{P} : (\exists p \in P)(\forall A \in \mathcal{A}, X \in \mathcal{X})\left(q(A,X) = \frac{s(p,A)(A,X)}{\|\mathcal{A}\|}\right)\right\}$; different (I think more intuitive than Type 1) act pricing.

1. Each act $A \in \mathcal{A}$ has a characteristic gamble $g_A(\omega) = g_A\big((A',X)\big) = \mathbb{I}_A(\omega)u(A',X)$.
2. Link between D and P_{eff}: $g \in D$ iff $\left(\forall p \in P_{eff}\right)\big(E_p(g) > 0\big)$.
3. Read act preferences from gamble desirability: $A \succ B$ iff $g_A - g_B \in D$.
4. Act pricing: agent will pay ϵ for A iff $g_A - \mathbb{I}_A\epsilon \in D$; agent will sell A for ϵ iff $\mathbb{I}_A\epsilon - g_A \in D$.

Pros & Cons

Type 1	Type 2
+ *D* directly represents *P*.	- *D* is based on an "effective" (fake) credal set.
+/- Pricing *looks* "familiar", but this is misleading (a constant gamble ϵ has a different meaning on a combined act-state space than in the normal context).	+ Pricing is more natural (I think, anyway).
- Weird "characteristic" gambles	+ Natural characteristic gambles
+/- Gamble values represent the "suppositional value" (what Andrew Bacon calls "actual value") of the act.	+/- Gamble values represent the utility achieved at each world consistent with the act; 0 when inconsistent.
Both representations exactly capture the act preferences and prices of the supervaluated SDT, just encoded differently.	
- Representation only exists for the special case of generalized imaging.	**+ Can be used to represent *any* SDT.**

Read the paper and supplement for more!

- Proofs (some in paper, others in the supplement).
- Coherence and consistency for SDGs on the total outcome space.
- Imprecision about the supposition rule itself.
- A simple numerical example; a slightly more complicated example is in the supplement.
- More discussion of the interpretation of the way pricing is represented in each bridge.
- Discussion of limitations/possible future work.
- And more …