University of BRISTOL

European Research Council

Established by the European Commission

Independent natural extension for choice functions

Arthur Van Camp,

Kevin Blackwell (live on your screen),

Jason Konek

4 Questions in 10 Minutes

1. How do sets of desirable gamble sets generalize sets of desirable gambles?
2. What does this generality add?
3. What is the independent natural extension?
4. What does our paper do and why should you care?

Sets of Desirable Gambles

- Consider the toss of a coin; we know that the coin will land on one of its faces.
- We have a set of possible outcomes: $\mathcal{X} = \{\text{Heads, Tails}\}$.
- In general, we assume that the set of outcomes forms a finite partition.
- There are continuum-many wagers on the outcomes that someone could offer us: $g(\mathrm{H}), g(\mathrm{T}) \in \mathbb{R}$.
- We call the space of all such possible wagers $\mathcal{L}(\mathcal{X})$.

Sets of Desirable Gambles

- If we assume that an agent is given decision problems where their only choice is to accept the gamble or decline it (maintain their present level of wealth), given assumptions about the decision rule the agent uses, we can interpret the subset $D \subseteq \mathcal{L}(\mathcal{X})$ that the agent is willing to accept as representing beliefs about whether the coin will land Heads or Tails.
- These beliefs are closely related to sets of probability functions, but not quite equivalent.
- Assume $g \in D \leftrightarrow (\forall p \in \mathcal{M}) E_p(g) > 0$, "Unanimous Positive Expectation," where $\mathcal{M}$ is a set of probability functions defined on $\mathcal{X}$.
- Then every set of probability functions $\mathcal{M}$ determines a unique set of desirable gambles D, but the converse is not true. [ItIP, 1.6.2]

Sets of Desirable Gambles

- Moving from a set of desirable gambles D to a credal set $\mathcal{M}_D$, there is information loss in two directions:

1. There are Ds that are not representable by any set of real-valued probability functions. (E.g., infinitesimal biases.) (But they *almost* are.)
2. For any credal set $\mathcal{M}$, $\mathcal{M}$'s convex hull yields the same set of desirable gambles. So D determines the extreme points of $\mathcal{M}_D$, but leaves the interior unspecified.

Sets of Desirable Gamble Sets

- Instead of assuming that an agent is posed binary decision problems, we might instead assume that they could be offered any finite non-empty option set consisting of gambles from $\mathcal{L}(\mathcal{X})$: $A \in \mathcal{Q}(\mathcal{L}(\mathcal{X}))$ iff A is finite and nonempty and $A \in \mathcal{P}(\mathcal{L}(\mathcal{X}))$.
- Which $A \in \mathcal{Q}(\mathcal{L}(\mathcal{X}))$ contain at least one gamble that the agent prefers to the status quo?
- Answer: the agent's set of desirable gamble sets! $K \subseteq \mathcal{Q}(\mathcal{L}(\mathcal{X}))$.

Sets of Desirable Gamble Sets

- Sets of desirable gamble sets are equivalent to the framework of choice functions, as introduced by Seidenfeld, Schervish, and Kadane [CCFuU]. [ADBAfCCF, Section 4]
- They generalize sets of desirable gambles in two important ways:

1. SDGS are decision-theoretically more informative than sets of desirable gambles.
2. SDGS allow us more (operationally) distinct representations of an agent's beliefs.

More Informative about Decision Rules

- Suppose we know that an agent's beliefs about the coin toss are well-represented by some credal set $\mathcal{M}$.
- Then you watch a bunch of bookies go up to the agent, each offering some gamble $g \in \mathcal{L}(\mathcal{X})$.
- You observe that the agent accepts all and only the gambles that have Unanimously Positive Expectation
  – $g: (\forall p \in \mathcal{M}) E_p(g) > 0$.
- Now suppose that there will be a second wave of bookies; this time, the bookies will offer arbitrary finite option sets $A \in \mathcal{Q}\big(\mathcal{L}(\mathcal{X})\big)$.
- Question: do you know how the agent will choose in this second wave?

More Informative About Decision Rules

- Answer: No! (Maybe this was obvious from the titles of the slides.)
- The agent's decision rule for binary choices that we observed in the first wave underdetermines their decision rule in the more general case.
- Many decision rules that make distinct judgments in the more general case collapse to the observed rule in the binary case.
- In particular, two of the most-discussed decision rules:
  - Maximality: $f \in C(A) \leftrightarrow (\forall g \in A)(\exists p \in \mathcal{M})\left(E_p(f) \geq E_p(g)\right)$
  - E-Admissibility: $f \in C(A) \leftrightarrow (\exists p \in \mathcal{M})(\forall g \in A)\left(E_p(f) \geq E_p(g)\right)$ both collapse to:
  - $0 \in C(\{0, g\}) \leftrightarrow (\exists p \in \mathcal{M})\left(0 \geq E_p(g)\right)$.
  - Then the binary sets of desirable gamble sets observed in the first wave are:

  $\{0, g\} \in K_{bin} \leftrightarrow 0 \notin C(\{0, g\}) \leftrightarrow \neg(\exists p \in \mathcal{M})\left(0 \geq E_p(g)\right)$
  $\leftrightarrow (\forall p \in M)\left(E_p(g) > 0\right)$
  - Which is exactly what you get from Unanimous Positive Expectation.

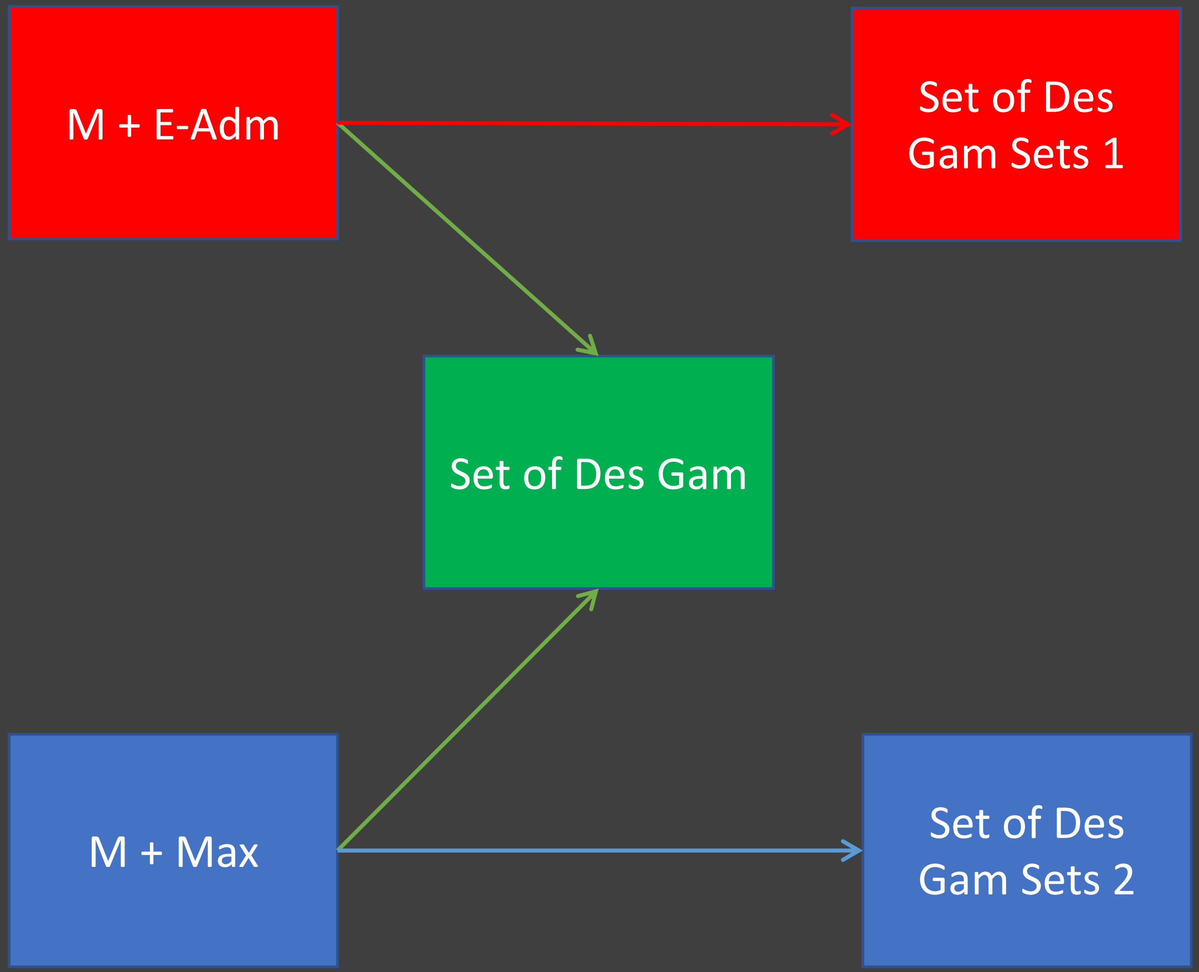

More (Operational) Representations of Belief

A is certain that a coin is either double-headed or double-tailed, but has no idea which; it’s natural to want to model A as represented by the credal set $\{(1,0),(0,1)\}$ (notation: (p_H, p_T)).

B is maximally imprecise, or vacuous, about the bias of the coin. B is represented by the set of all probability functions defined on the outcomes of the coin toss.

More (Operational) Representations of Belief

- Both A and B's credal sets generate the same set of desirable gambles, namely the vacuous model:

$\mathcal{L}_{>0} = \{g \in \mathcal{L}(\mathcal{X}): g(H), g(T) \geq 0 \text{ and } (g(H) > 0 \text{ or } g(T) > 0)\}$.

- But, choosing via E-Admissibility, they generate different sets of desirable gamble sets.
    1. B's SDGS is also the vacuous model $K_V = \{A \in \mathcal{Q}(\mathcal{X}): (\exists f \in A) f \in \mathcal{L}_{>0}\}$.
    2. A's SDGS is not just the vacuous model. It also includes sets of gambles like $\{\langle 1, -100\rangle, \langle -100, 1\rangle\}$.

What is Independence?

Our paper includes two independence concepts, each of which we think of as intuitively compelling.

1. Epistemic independence: suppose I'm going to learn some proposition $(E_Y \subset \mathcal{Y})$ about the value of Y and I want to know: is there anything I could learn about the value of Y that should change my beliefs about X? If the answer is no, Y is *epistemically irrelevant* to X. *Epistemic independence* is the symmetric version of this kind of irrelevance: X is irrelevant to Y, Y is irrelevant to X.

- The primary notion of independence we work with in our paper is a formalization of this concept in the context of sets of desirable gamble sets.

What is Independence?

2. Practical independence: rather than taking it as given that you *do* learn about one of the variables, now think of learning as a choice that has a cost. I could, right now, make some decision that depends on my beliefs about X or I could pay (literal money, time, effort, anything of value) to learn about Y. Should I choose to learn before deciding? If there is nothing of value that I am willing to exchange to learn about Y before deciding about X, I judge Y to be *practically irrelevant* to X. *Practical independence* is the symmetric version of this irrelevance concept.

- Teddy Seidenfeld has proposed a formalization of this sense of independence [CCFuU], which has also been studied in great detail by Jasper De Bock and Gert de Cooman [OaNoIPbTS].

“The” Independent Natural Extension

- Suppose we have two partial assessments: A_X and A_Y which reflect, respectively, beliefs about the variables X and Y.
- In general terms, the idea behind the independent natural extension A of these assessments is that:
  - Marginalized to one of X or Y, A returns the partial assessment A_X or A_Y, as appropriate.
  - A represents X as irrelevant to Y and Y as irrelevant to X.
  - A is the least-informative/least-committal coherent assessment (of the same kind as A_X and A_Y) which does both of the above.

“The” Independent Natural Extension

- The intuitive idea is that, given a class of models of beliefs with associated coherence conditions and a formalized notion of what it means for two variables to be mutually irrelevant, the independent natural extension represents all of the beliefs that an agent *must also commit to* in order to coherently believe the two partial assessments A_X and A_Y while holding that X and Y are mutually irrelevant.
- Put another way: it collects all and only the beliefs common to every model that coherently judges X and Y to be independent while satisfying the partial assessments A_X and A_Y.
- “The” is in scare quotes, because there are two degrees of freedom:
    1. Choice of model class / associated notion of coherence.
    2. Choice of independence concept.

What Our Paper Does

Our paper has two distinct parts.

1. In the first, we build on the work of De Cooman and Miranda, who identified and proved the form that the *epistemically* independent natural extension takes for *sets of desirable gambles*. We extend this same formalization of epistemic independence to sets of desirable gamble sets, identifying and proving the form that the independent natural extension takes in this more general model class [IINEfSoDG].
   - Our poster focuses on one nice example of the kind of extra generality that moving to sets of desirable gamble sets buys: we show how to independently combine partial assessments on two different variables generated, respectively, by maximality and E-Admissibility. To the best of our knowledge, this is entirely novel.

What Our Paper Does

2. In the second, we hold fixed the model class (sets of desirable gamble sets) and contrast two different notions of independence that we find intuitively compelling and important: the *epistemic* independence from the first part with Teddy Seidenfeld's *practical* independence notion.
   - In their study of this S-independence, de Bock and de Cooman have already shown that epistemic independence does not entail S-independence.
   - We provide an example (also on the poster!) which demonstrates the converse direction: S-independence does not entail epistemic independence.
   - It is perhaps somewhat surprising that neither of these independence notions entails the other.
   - (Although, in light of investigations of violations of Good's Value of Information theorem in IP, perhaps not so surprising? E.g., [CFEBB?VoIftIP])

Secret Bonus Question: Should you come to our poster session?

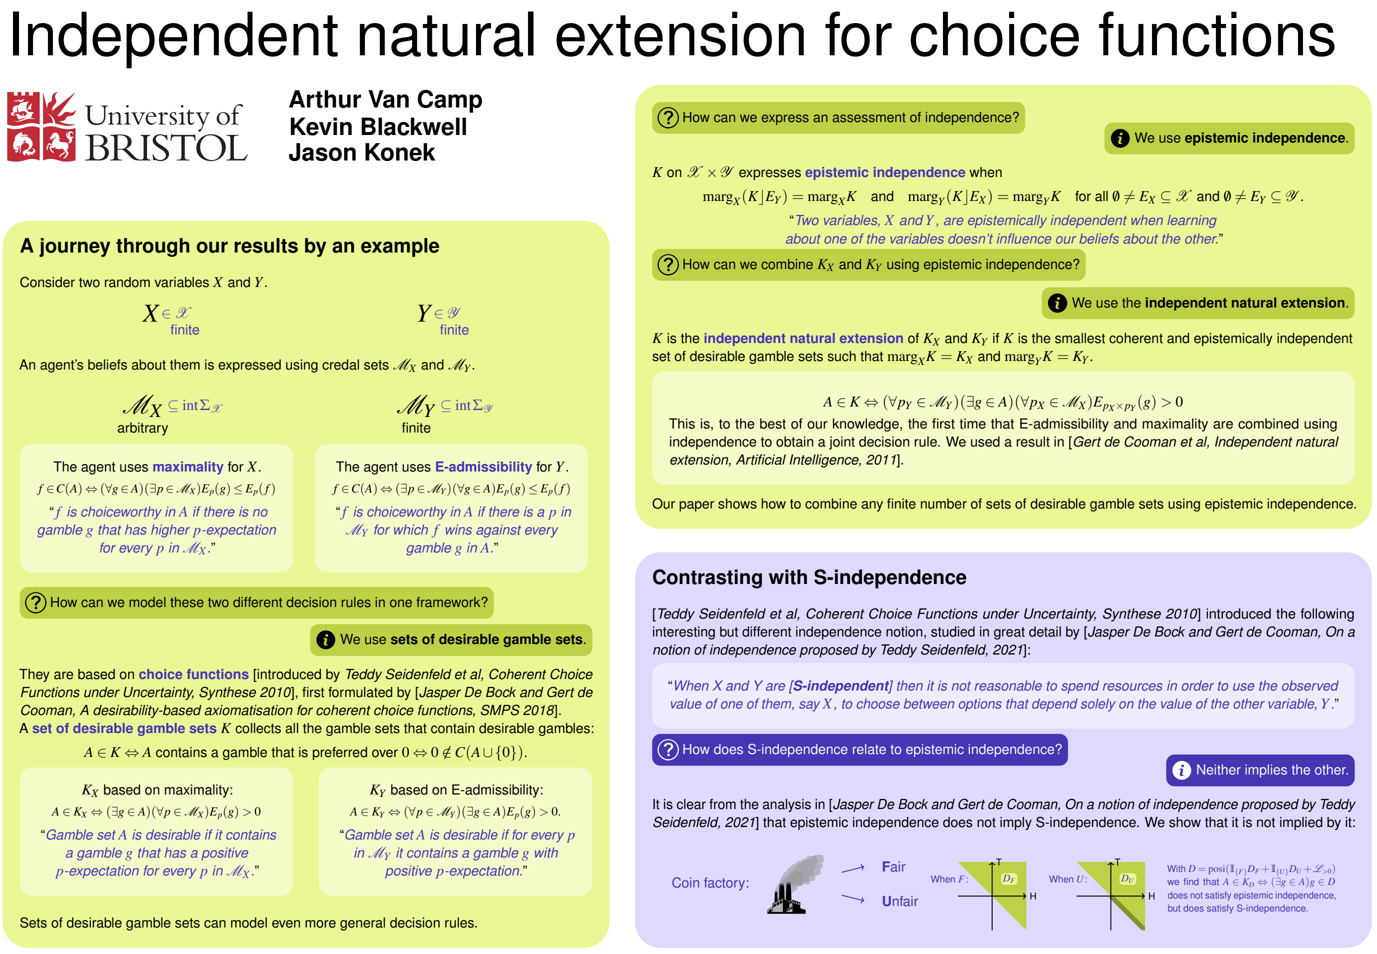

Answer: left as an exercise to the audience.

References (in order of appearance in slides)

[ItIP] T. Augustin, F. P. Coolen, G. de Cooman, and M. C. Troffaes, eds. 2014. *Introduction to Imprecise Probabilities*. Hoboken, NJ: Wiley.

[CCFuU] Teddy Seidenfeld, Mark J. Schervish, and Joseph B. Kadane. 2010. “Coherent choice functions under uncertainty”. *Synthese*, 172(1):157–176. doi: 10.1007/ s11229-009-9470-7.

[ADBAfCCF] Jasper De Bock and Gert de Cooman. 2018. “A desirability-based axiomatisation for coherent choice functions”. In S. Destercke, T. Denoeux, M.A. Gil, P. Grzer- gorzewski, and O. Hryniewicz, eds. *Uncertainty Modelling in Data Science: Proceedings of SMPS 2018:* 46–53.

[OaNoIPbTS] Jasper De Bock and Gert de Cooman. 2021. “On a notion of independence proposed by Teddy Seidenfeld”. URL https://arxiv.org/abs/2102. 10342.

[IINEfSoDG] Gert de Cooman and Enrique Miranda. 2012. “Irrelevant and independent natural extension for sets of desirable gambles”. *Journal of Artificial Intelligence Research* 45:601–640.

[CFEBB?VoIftIP] Seamus Bradley and Katie Steele. 2015. ”Can Free Evidence Be Bad? Value of Information for the Imprecise Probabilist”. *Philosophy of Science* 83(1): 1-28.