Formalising two conceptions of probability

Probability's "squirming duality"*

- Subjective or objective?
- Degrees of belief or (limiting) frequency?
- We'll dodge the philosophical controversies . . .
- . . . and formalise both views

*Hacking

A first-order logic of probability

- Taken straight from Halpern
- Skipping the tricky bits
- *Reasoning with Uncertainty* (MIT Press) is recommended reading

The basic dichotomy formalised

- "The probability that a randomly chosen bird will fly is greater than 0.9" — objective

- $w_x(\mathit{flies}(x)) > 0.9$

- "The probability that *tweety* (a particular bird) flies is greater than than 0.9" — subjective

- $w(\mathit{flies}(tweety)) > 0.9$

‘Sampling’ probabilities - syntax

- Start with a first-order language (with a collection of function and predicate symbols Φ) which allows us to say things like $\mathit{flies}(\mathit{tweety})$

- But for any first-order formula ϕ we extend the language to include *probability terms* $w_{\mathbf{x}}(\phi)$: the probability that a randomly chosen $\mathbf{x}$ satisfies ϕ.

- $w_{x,y} son(x, y) = 0.2$

- We also include terms to represent numerical probabilities.

‘Sampling’ Probabilities - Probabilities on the domain

A *type-1 probability structure* is a tuple (D, π, μ) where

- D is a domain
- π maps function and predicate symbols in Φ to to functions and predicate symbols of the right arity over D
- μ is a discrete probability function on D.

So it's a standard first-order structure (D, π) ...
...with a probability distribution μ tacked on.

A world where tweety flies - and most other things too

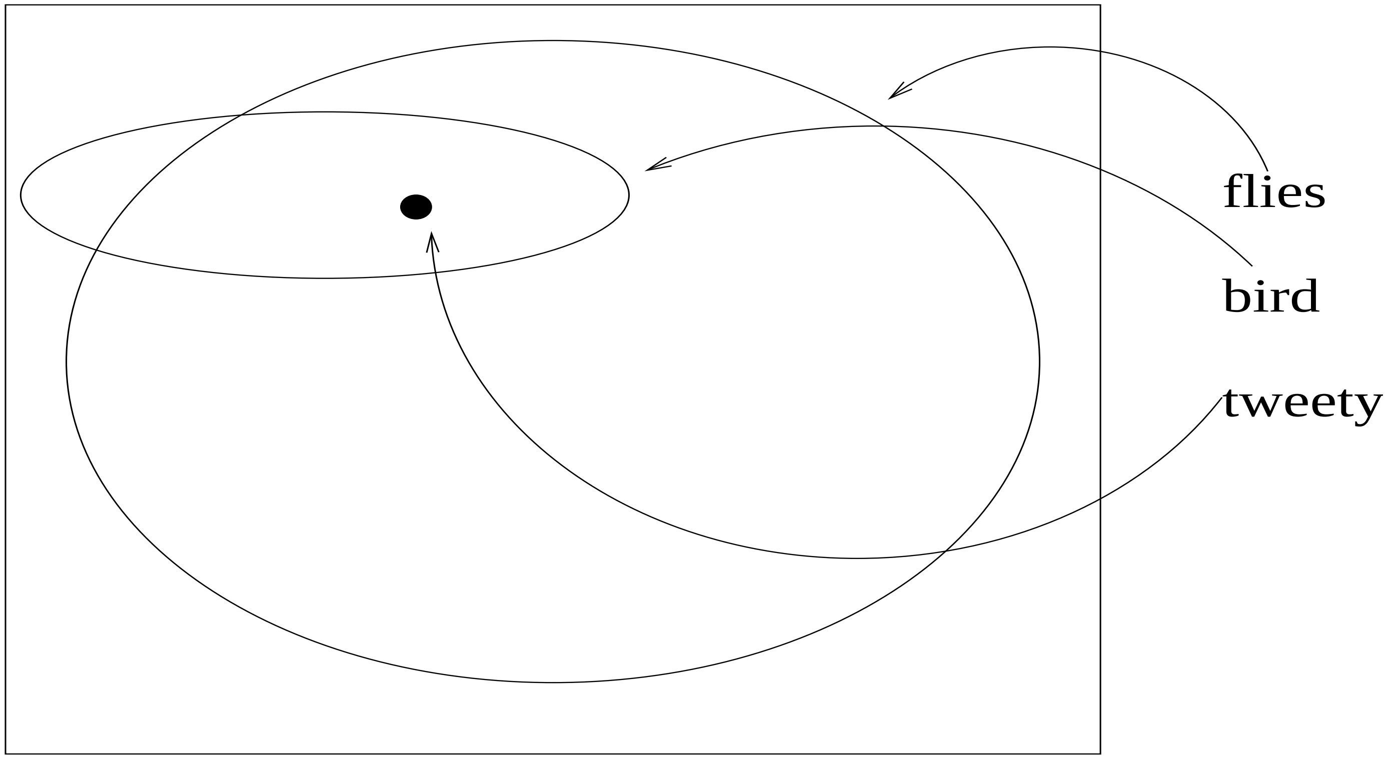

Semantics for probabilities on the domain

Let $M = (D, \pi, \mu)$ be a type-1 probability structure, then

- $M \models \mathit{flies}(\mathit{tweety})$ iff $\pi(\mathit{tweety}) \in \pi(\mathit{flies})$

- $M \models w_x(\mathit{flies}(x)) > 0.9$ iff $\mu(\{d : d \in \pi(\mathit{flies})\} > 0.9$

So $\forall x \mathit{flies}(x) \models w_x(\mathit{flies}(x)) = 1$
But $\mathit{flies}(\mathit{tweety}) \not\models w_x(\mathit{flies}(x)) > 0$

‘Degree-of-belief’ probabilities - syntax

- Start with a first-order language (with a collection of function and predicate symbols Φ) which allows us to say things like $\mathit{flies}(\mathit{tweety})$

- But for any first-order formula ϕ we extend the language to include *probability terms* $w(\phi)$: the probability that ϕ is true.

- We also include terms to represent numerical probabilities.

- So we can say $w(\mathit{flies}(\mathit{tweety})) = 0.2$

Two possible worlds

In one world *tweety* flies, in the other he doesn't.

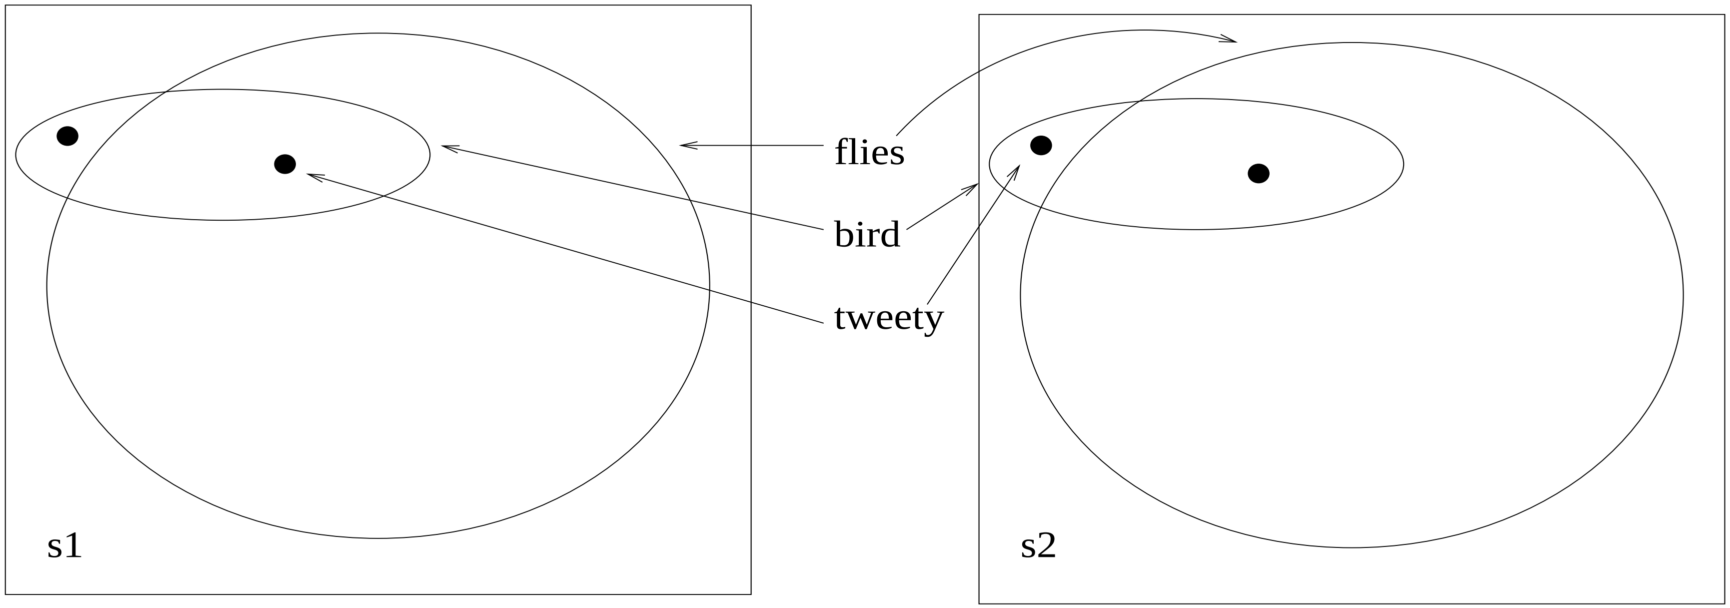

‘Degree-of-belief’ probabilities - probabilities on possible worlds

A *type-2 probability structure* is a tuple (D, S, π, μ) where

- D is a domain
- S is a set of possible worlds
- For each $s \in S$: $\pi(s)$ maps function and predicate symbols in Φ to to functions and predicate symbols of the right arity over D
- μ is a discrete probability function on S.

Two possible worlds - revisited

Suppose $S = \{s_1, s_2\}$, where $\pi(s_1)(tweety) \neq \pi(s_2)(tweety)$

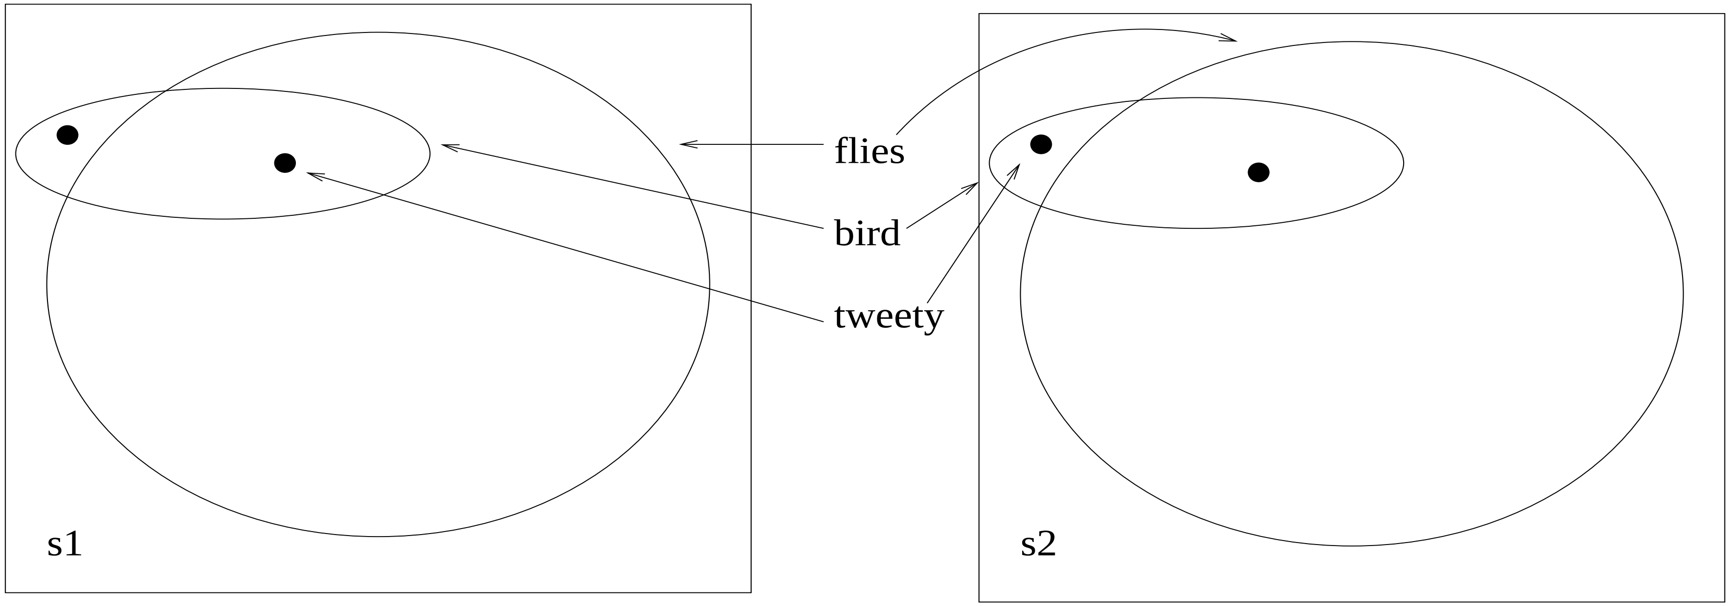

We might have, say $\mu(s_1) = 0.2$ and $\mu(s_2) = 0.8$

Semantics

Let M be a type-2 probability structure and $s \in S$ a world, then

- $(M, s) \models \mathit{flies}(\mathit{tweety})$ iff $\pi(s)(\mathit{tweety}) \in \pi(s)(\mathit{flies})$
- $(M, s) \models \mathit{flies}(\mathit{tweety}) = 0.2$ iff
  $\mu\{s' \in S : (M, s') \models \mathit{flies}(\mathit{tweety})\} = 0.2$

$\mathit{flies}(\mathit{tweety}) \not\models w(\mathit{flies}(\mathit{tweety})) = 1$