An austere physical description of the world starts from physical fields and forces distributed in spacetime. These austere ingredients sometimes make up solid material things, and sometimes material things form systems with the capacity to change the world around them in predictable ways. Parts of a material system have *functions*, in virtue of the way that their properties explain how the overall system can do what it does. In some cases, the systems reproduce themselves, and the functional properties of parts are selected-for. This can ground the truth that a part is *supposed* to function a certain way, even though it is currently not functioning that way.

We have moved from the austere physical description of the world to a comparatively rich one: a world of material things with a naturalized teleology. But the world in which we live has many further features. For each such feature we can ask: can they be built out of the materials just mentioned? If not, what further materials are required?

In *The Metaphysics of Representation* takes on one task of this kind. Our world is one in which certain systems perceive things, believe that things are one way, desire them to be another way, form intentions and communicate with one another. I want to understand how these kind of representational facts are grounded in more basic materials.

Some kinds of representation are close to the surface of a world under the teleological physical description. Your thermostat has the capacity to keep your room at constant temperature, a capacity explained in part by the temperature sensor within it. The sensor, say, is a bimetallic strip, which deforms in specific ways depending on the ambient temperature. The thermostat will be functioning correctly only if the various states of the sensor are being produced by specific room-temperatures. The deformation of the sensor *represents* a specific room temperature range in virtue of it having the function (within the thermostat-system) to be produced by that specific state of affairs. I join teleosemanticists in saying (roughly) that certain states within biological systems—in particular, sensory-perceptual states within the human biological system—have representational content in this way. I give a parallel account of the representational content of intentional-motor states.

This simple story can't be applied directly to representational states such as beliefs or desires (or to the concepts within them). A normal human has plenty of beliefs that are not produced by the state of the world they represent, despite the fact that they're functioning perfectly well. Sally believes on good evidence that the back wall of her garden will be standing tomorrow; but in fact, it falls. Harry has plenty of desires that do not bring about the state of the world they depict: he wants world peace, but is disappointed. The failure of his desire to bring about its satisfaction is due to his limited influence, not any cognitive malfunction.

The teleosemanticist's bet is that we can find a more sophisticated version of the teleological reductive story, which will cover not only perception (and intention), but all the other representational states we are interested in. By contrast, I stick with the simple story, and use the representational resources already grounded as materials from which to build the further representational facts about belief and desire.

So we've earned the right to describe our perceptual and intentional-motor subsystems in representational terms, and now face the challenge to pinpoint the content of the internal states that mediate between these perceptions and intentions. I help myself to a new ingredient, over and above the distribution of facts about physical fields and forces that lie at the basis of what we've used so far. The new basic ingredient are normative facts about when beliefs would be justified, given evidence, and when actions are justified, given beliefs and desires. The hypothesis is this: among all the possible interpretations of those internal states, the correct one is that which makes the agent most rational. An inner belief or desire state has representational content that p in virtue of the correct interpretation (so understood) assigning p to that state. This story is one precisification of "radical interpretation".

Feed the physical facts about a human being into this two-part story, and the relevant facts about rationality required will look familiar: our typical theories of epistemic and practical justification are

after all, designed for the human case. Presuppositions about our contingent cognitive setup are often present in such theories. For example: it's plausible that our folk theory of rational deduction presupposes that our beliefs have language-like structure. By making contingent assumptions about human cognitive architecture and human rationality, we can (conditionally) predict and explain local patterns in the way our concepts pick out features of the world: why moral concepts are referentially stable; why natural properties are reference magnetic; how we manage to quantify over absolutely everything. Analogous predictions about an alien psychology should be possible, but would require a specification of alien psychological stucture and alien rationality (or an abstract, presupposition-free theory of rationality) which we don't have ready to hand.

You and I possess a shared public language in which we express and communicate our thoughts; other actual and possible thinkers do not. In the story I tell, when two creatures share a public language, that matters to how we interpret their beliefs and desires, since it links the interpretations of their respective concepts paired with one and the same public language expression. The presence of public language inflects the metasemantics of thought, but I say that this is consistent with the view that the content of language is always inherited (via linguistic conventions) from the content of thought.

The upshot is a three-layer story about representation, each layer building on the previous: teleosemantically grounded perceptual and intentional representation; interpretationistically grounded belief and desire representation; conventionally grounded linguistic representation.