It is tempting to think that people ***should not*** believe some of the things they believe. Think of your favorite example, e.g., the vaccine-denier who has plenty of evidence against her belief that the polio vaccine is dangerous, or closer to home, a colleague who has plenty of evidence that he doesn't do his share of departmental administrative duties and yet thinks he is above average in that respect. But sometimes our beliefs don't fit our evidence in ways it would take someone brilliant to recognize, e.g., think of an ordinary person (not a mathematician) who sees there are 25 people in the room, and after reflecting a bit, concludes that it's less than 50% likely that two or more of them share a birthday. Suppose this person is aware of the probability axioms along with basic facts such as people have just one birthday and there are 365 days in a year. The axioms plus these facts and a little logic entail that his belief is false. But should this person not believe this? Should he give up his belief? We would like to answer: *no, that's asking too much of him*. But what is the difference between this case and the other ones ? I explain the difference in terms of abilities or capacities to appreciate the force of specific bodies of evidence.

Giving abilities this role, in turn, has significant theoretical consequences. I discuss three. First, it shows that a popular view called "evidentialism" needs amendment. Second, it supports a *path-first* view about belief-revision: we should give up a belief only if we're well-positioned to take a *path* from our evidence to giving it up. Such paths are ways of "following one's evidence." Third, giving abilities this role gives us resources for building an argument that people ***should believe*** certain things, contrary to the claims of purely "permissive" epistemological theories.