

9

Moral Information

Introduction

In this chapter, we turn to a final implication of moral uncertainty: the reason, in terms of expected choice-worthiness, it gives to gain new moral information. In what follows, we introduce a framework for understanding this.

In section I, we explain how we should assess the expected utility of new empirical information, and how we could use an analogous framework to work out the expected choice-worthiness of new moral information. In section II, we apply this framework to two examples: the choice of how a large foundation should spend its resources, and the choice of career for an individual. In section III, we consider to what extent the lessons from this framework change when we consider ‘imperfect’ information.

Before we begin, we should highlight that we use the unusual term ‘moral information’. We use this term in the hope of remaining almost entirely non-committal on the issues of moral epistemology and moral metaphysics: as we understand it, something is a piece of moral information iff coming to possess it should, epistemically, make one alter one’s beliefs or one’s degrees of belief about at least one fundamental moral proposition. So, the term ‘moral information’ could apply to experiences, arguments, intuitions, or knowledge of moral facts themselves.

I. Assessing Gains in Information

In this section, we’ll explain how one should calculate the expected utility of gaining empirical information (understanding ‘utility’ as the numerical representation of the agent’s preference ordering). One can work out the expected utility of *perfect* information—that is, the expected utility of coming to know some particular proposition for certain—and the expected utility of *imperfect* information, which is the expected utility of improving one’s

evidence base but not coming to know any additional proposition for certain. We'll begin by discussing the simpler concept of the expected utility of perfect information, discussing imperfect information in section III. We'll illustrate the idea of perfect information with recourse to the following example,¹ before discussing the idea in general.

Salesman

Jonny sells ice-cream cones. He has fifty ice-cream cones and he makes \$1 profit for each ice-cream cone he sells. He has the option of reserving a market stall by the beach for a day for \$10. If the weather is sunny, then he will sell all his ice-cream cones; if it is raining, he will sell none. He has the option to purchase access to a new incredibly reliable meteorological service, which can tell him for certain whether it will rain or be sunny tomorrow. How much, rationally, should Jonny be willing to pay in order to know for certain how many ice-cream cones he'll be able to sell if he tries?²

According to the standard decision-theoretical analysis,³ he should answer this question as follows. First, he should work out how many ice-cream cones he expects to sell, given his current evidence. Let's suppose that he thinks there is a 50/50 chance of rain. Second, he should work out the expected utility of his options, given his current evidence. In this case, the expected utility of not-reserving a market stall is 0. The expected utility of reserving the stall is $0.5 \times (\$50 - \$10) + 0.5 \times -\$10 = \15 . The expected utility of reserving the stall is higher than the expected utility of not reserving the stall. So, given his current evidence, he should reserve the stall.

Third, he should work out the additional utility of gaining the new information. If he finds out that it will be sunny, then the additional information has no utility for him: he would not change his behavior with this new information, and so he would have made the same amount of money even without the new information. However, if he were to find out that it will rain, he would change his behavior: he would decide against reserving the stall. So, if it is the case that it will rain, the utility for Jonny of finding that out is \$10.

¹ This version is adapted from the 'newsboy' example given in Louis Eeckhoudt and Philippe Godfroid, 'Risk Aversion and the Value of Information', *The Journal of Economic Education*, vol. 31, no. 4 (Autumn 2000), p. 382–8.

² We'll also make some simplifying assumptions: that Jonny doesn't value his time at all, that this is a one-time opportunity, and that the value of additional dollars for Jonny is linear over this range.

³ See, for example, Howard Raiffa, *Decision Analysis: Introductory Lectures on Choices Under Uncertainty*, Reading, MA: Addison-Wesley, 1968.

Now, Jonny thinks there is a 50% chance that he will find out that it will be sunny (which would have no utility for him), and a 50% chance that he will find out that it will rain (which would be worth \$10). So the expected utility of gaining that new piece of information is $0.5 \times \$0 + 0.5 \times \$10 = \$5$. This gives the amount up to which he should be willing to pay for the meteorological report.

In general, the expected utility of gaining new information is given by the expected utility of one's decision given that new information (in this case, \$20 ($0.5 \times (\$50 - \$10) + 0.5 \times \0)) minus the expected utility of one's decision without that information (in this case, \$15 (calculation given above)).

When dealing with the expected utility of information, there are some important points to note. First, as one might have noticed from the above, on our analysis, gaining new information has positive expected utility *only* if there is some chance that one will change one's behavior. If Jonny thinks that he would sell forty ice-cream cones even if it were raining, then there is no expected utility for him in gaining additional information, because he would reserve the stall either way. Similarly, if Jonny knows that he is very lazy, and will fail to reserve the stall no matter how rational it is for him to do so, then, again, gaining new information will have no expected utility for him. In reality, factors such as peace of mind can make it rational to gain new evidence even if one will not change one's behavior. But for simplicity, we leave these details to the side.

Second, note that the expected utility of information is very different from how much one actually has to pay for that information. Perhaps Jonny could find out whether it will rain tomorrow simply by checking online, costing him nothing. In which case, he simply had a bargain—but the amount he had to pay does not change the fact that the information had an expected utility of \$5 (and that, if he had no better option, he should have been willing to pay up to \$5 to receive it).

Third, the higher stakes a decision is, the greater the expected utility of information. To illustrate, suppose in the case above that we multiplied all the monetary values by 10: each ice-cream cone sells for \$10, but Jonny has to pay \$100 in order to reserve the stall. In which case, the expected utility of information for Jonny would have the same proportional change, increasing to \$50.

The above method for calculating the expected utility of additional information is intuitively appealing and widely accepted within decision analysis. But, to our knowledge, it has only ever been used to work out the expected utility of gaining new empirical information: that is, information

about how the world will pan out. One unique evaluation of all possible states of the world is always presupposed.

Improving our epistemic state with respect to the moral facts is something we can do, and something that could potentially change which actions we take and believe it's appropriate to take. If our argument in this book so far is correct, then it seems plausible that we should assess the expected choice-worthiness of gaining more information about moral facts in just the same way that we should assess the expected utility of gaining more information about empirical facts.

Given this, the expected choice-worthiness of gaining new information is given by the expected choice-worthiness of one's decision given that new information minus the expected choice-worthiness of one's decision without that information.⁴

In what follows, we'll give two examples to illustrate some applications of this analysis to moral information.

II. Two Examples

A Philanthropic Foundation

Our first example provides the simplest illustration of the expected choice-worthiness of moral information.⁵ Let us suppose that the leader of a major philanthropic foundation is deciding how to allocate \$10 million of her resources. She is deciding between two possible grants. The first would go to the Against Malaria Foundation (AMF), which she believes would provide, on average, one extra year of very high-quality life (one 'Quality Adjusted Life Year' or 'QALY') to the extreme poor for every \$100 it receives.⁶ The

⁴ Formally, for some piece of information I and different ways j that the information could turn out, the expected choice-worthiness of gaining that piece of information I is: $\sum_{j=1}^n [C(I_j)(\max_A \sum_{i=1}^n C(T_i | I_j) CW_i(A) - \max_A \sum_{i=1}^n C(T_i) CW_i(A))]$. Note that this assumes a perfectly rational and enkratic agent.

⁵ We give this example as a hypothetical, but it is relevant to real-life cases, in particular to the foundation Good Ventures, advised by the Open Philanthropy Project and its sister organization GiveWell. For a discussion of their uncertainty concerning different 'worldviews' (which includes moral uncertainty), see Holden Karnofsky, 'Worldview Diversification', *Open Philanthropy Project*, 13 December 2016, <http://www.openphilanthropy.org/blog/worldview-diversification>; Holden Karnofsky, 'Update on Cause Prioritization at Open Philanthropy', *Open Philanthropy Project*, 26 January 2018, <https://www.openphilanthropy.org/blog/update-cause-prioritization-open-philanthropy>.

⁶ GiveWell, 'Mass Distribution of Long-Lasting Insecticide-Treated Nets (LLINs)'.

second would be to support corporate cage-free egg campaigns. She believes that every \$100 received by this campaign will ultimately cause farms to raise 3,800 laying hens (each of which live for approximately one year⁷) in a cage-free environment rather than a cage.⁸ For simplicity, we'll stipulate that the foundation leader is certain of consequentialism.

Let's further suppose that the leader of this foundation is certain of the moral value of one QALY (for a human), so the current expected choice-worthiness of the grant to the Against Malaria Foundation is 100,000 QALYs. However, she is extremely uncertain about the value of improving conditions in factory farms: she is 99% certain that there is no value to ensuring that chickens live in a cage-free environment; she has 1% credence that the value of ensuring that a single hen is brought up in a cage-free environment (rather than that a different hen is brought up in a caged environment) is 1/100th of the value of a QALY. She believes that, across these two moral views, the value of one QALY stays constant.

Given this, the grant to Against Malaria Foundation has an expected choice-worthiness of 100,000 human QALYs,⁹ whereas the grant to the cage-free egg campaigns has an expected choice-worthiness equivalent to only 38,000 human QALYs, so the best decision, given the credences she has, is to give the grant to the Against Malaria Foundation.

Now, suppose that the decision-maker has the option of gaining perfect information about the relative value of improving the conditions of layer hens versus providing one QALY. What's the expected choice-worthiness of this information? We can work this out using the framework given above. She should think that there's a 99% chance of finding out that the cage-free reforms are of no value, so gaining this information is 99% likely not to change her behavior, and therefore have no value (at least, within the context of this decision). But she should think that there's a 1% chance that she will learn that the cage-free campaigns are of value: if this happened,

⁷ 'The Life of Laying Hens,' *Compassion in World Farming*, March 2012, <https://www.ciwf.org.uk/media/5235024/The-life-of-laying-hens.pdf>.

⁸ Lewis Bollard, 'Initial Grants to Support Corporate Cage-free Reforms,' *Open Philanthropy Project*, 31 March 2016, <http://www.openphilanthropy.org/blog/initial-grants-support-corporate-cage-free-reforms>. Of course, the hens that are raised in a cage-free environment are not the same hens as those that would have been raised in a cage. And, though it is not realistic to suppose that farms will raise exactly as many chickens when they are raised in a cage-free environment, we make this assumption for simplicity.

⁹ Strictly speaking, QALYs are a unit of goodness rather than choice-worthiness, so a more accurate (but more cumbersome) way of saying the above is that there is an expected choice-worthiness equivalent to the choice-worthiness that theories ascribe to producing 100,000 QALYs for humans.

then the expected choice-worthiness of the grant to the cage-free campaigns would become equivalent in size to providing 3.8 million human QALYs. The additional benefit she would produce given this outcome would therefore be of equivalent value to providing 3.7 million QALYs. Multiplying the value of this outcome by its probability of 1% gives us the expected choice-worthiness of gaining the information, which is equivalent to providing 37,000 QALYs. The cost to provide one QALY via a donation to AMF is \$100. So she should be willing to pay up to \$3.7 million (that is, $37,000 \times \$100$) in order to gain this information before making her decision about where to spend the \$10 million.

In the above calculation, her starting budget was not relevant. It turned out that she should spend \$3.7 million to help direct her \$10 million. This means spending \$3.7 million out of a total spend of \$13.7 million, which is 27%. Given the other details, these percentages stay the same, so regardless of her budget size she should be willing to spend about 27% of her budget in order to know how she ought to spend the remaining 73%. Thus, if her total budget were fixed at \$10 million, then she should be willing to spend \$2.7 million in order to find out how to spend the remaining \$7.3 million.

The above example is highly idealized, with invented numbers for the moral views and their credences, as well as a convenient restriction to just two possibilities for the value of a year of a chicken's life. But it's not completely unrealistic: we deliberately chose empirically accurate numbers, and we chose credences in moral views that could have (in a very simplified form) represented the views of the leadership of the Open Philanthropy Project at one time. The example shows, therefore, that it's at least possible for the expected choice-worthiness of moral information to be very high, such that a significant proportion of one's resources should be spent on gaining new moral information. (We'll discuss later to what extent the fact that moral information is almost always 'imperfect information' changes things.) In general, because information brings about a proportional change in the expected choice-worthiness of the options under consideration, if you're dealing with extremely high-stakes issues, then the expected choice-worthiness of gaining new information becomes extremely high as well.

This is notable given that philanthropists (and other similar actors, like governments) almost never spend resources on gaining new moral evidence. The typical method for a foundation, for example, is to pick a cause area to focus on (such as education, or climate change), and then use their resources to try to optimize within that cause area. However, they typically

spend very few resources to improve their answer to the question of what cause area it is most important for them to spend their resources on, even though answering that question will necessarily require making ethical judgements.¹⁰

Career Choice

As well as spending money to gain new moral information, one can also spend time gaining new moral information. This is relevant, for example, to the question of how much time young people should be willing to spend studying ethics before choosing which career to pursue. Again, we'll give an idealized example to illustrate. Consider Sophie. She comes from a poor family in the UK, but is very bright and hardworking, and won a scholarship to a top university. She's undecided about what career to pursue. She could become an NGO worker, and through that save the lives of one hundred people in developing countries, but it would mean that she could not give back to her family at all. Or she could become a lawyer: this would not benefit those in developing countries at all, but would mean that she could pay for health insurance and better living conditions for her extended family, improving the overall lives of each of twenty-two of her family members by 30%. She therefore realizes that she can benefit those in developing countries much more than she can benefit her family. But she isn't sure how to weigh those respective benefits. We'll assume, for simplicity, that she's certain in consequentialism. She's 95% confident that it's one hundred times more important to benefit her family, but has 5% credence remaining that it's just as important to benefit those in developing countries as it is to benefit her family, and that the moral value of benefiting her family stays constant across these two possible moral views. Given this, how much time should she be willing to spend studying ethics if doing so could give her perfect information about how to value benefits to her family compared with benefits to those in the developing world?

In what follows, we'll stipulate that saving one life in a developing country, according to the partial view, is worth 1 unit of value, and that benefiting someone's life by 30% provides 0.3 times as much benefit as

¹⁰ A notable exception is Good Ventures and the Open Philanthropy Project, which undertake significant investigation to try to make cross-cause comparisons.

saving someone's life. Given her current beliefs, it is appropriate for Sophie to choose to become a lawyer: the expected choice-worthiness of doing so is $0.3 \times 22 \times 100 = 660$ units of value, whereas the expected choice-worthiness of becoming a NGO worker is $0.95 \times 100 + 0.05 \times 100 \times 100 = 595$ units of value. But she also has the option of getting more moral information: she could take several years out before university in order to study moral philosophy. How many years should she be willing to spend studying in order to get perfect information about how to weigh benefits to her family against benefits to those in the developing world?

In this example, she should think it 95% likely that she wouldn't change her decision, as this is her credence that the partial view will turn out to be correct. But she should think it 5% likely that she would change her decision (as a result of discovering that she should be impartial between distant strangers and her family) and that by choosing to become the NGO worker she would increase the value of her career (by $100 \times 100 - 0.3 \times 22 \times 100 = 9340$). So the expected choice-worthiness of this information is $0.05 \times 9340 = 467$. So she should be willing to lose out on 467 units of value in order to gain perfect information about how to spend her forty-year career. Assuming that the benefit to her family were spread evenly over a 40-year career, she produces $0.3 \times 22 \times 100 / 40 = 16.5$ units per year. So she should be willing to spend $28.3 / (28.3 + 40) = 41.4\%$ of her time to gain perfect information about how to spend the remaining 58.6%. So, if she only had those forty years to spend, she should be willing to spend a little over sixteen of them studying ethics if this would give her perfect information about what she should do with the remainder of her career.

Like the previous example, this example was illustrated with invented credences, out of necessity. But it at least shows that the expected choice-worthiness of additional moral information *can* be high. And the thought, at least, that it could be worth anyone spending a significant proportion of their life studying ethics just so that they make a better decision at the end of that time is surprising. Indeed, for most non-philosophers, the thought that one should spend *any* time studying ethics before making major life decisions might be surprising.

Of course, in the above case the conclusion is not that Sophie actually should spend sixteen years studying ethics. Again, we need to distinguish the expected choice-worthiness of gaining moral information from the 'price' of that information—how much time it would actually take to get that information. Perhaps Sophie would learn most of what she needs to after only a few years of study. In which case it might no longer be worth

spending the remaining decade learning a bit more. But that does not diminish the expected choice-worthiness of those few years of study—it just means that, for those few years, she is getting a bargain, evaluatively speaking.

A second caveat, when it comes to how much time the typical person should spend studying ethics, is that the above assumption that the benefit Sophie would produce is linear over a forty-year career will likely often be inaccurate. It seems plausible that the benefit one produces in one's career increases dramatically over the course of one's life, as one gets promoted, and becomes more experienced and more influential. In which case, insofar as studying ethics pushes back one's career, thereby taking years off the end of one's career, the cost of studying ethics is higher than the above calculation would suggest. And one can lose career options by studying ethics for too long, providing another reason against too many years of study. Rather than sacrificing 41.4% of her time to gain perfect information, she should be willing to sacrifice enough time to reduce her future earnings/impact by 41.4%, which may be somewhat less.

But even despite these caveats, as with the previous case it seems plausible that the expected choice-worthiness of gaining new moral information is higher than one might expect. It seems perfectly plausible that being in a better epistemic state with respect to the moral facts can mean that one does ten times as much good in the rest of one's life as one would otherwise have done (e.g. perhaps one focuses on climate change mitigation rather than a domestic policy issue because one comes to believe that future people are much more important than one had thought). In which case, it would be worth spending half one's working life studying ethics in order to improve how one uses the remaining half—even if 80% of the value that one contributes to the world typically occurs in the latter half of one's career.

III. Imperfect Information

In the above examples, we assumed for ease of presentation that we'd be able to achieve certainty in the moral facts of the matter. But that's unrealistic: we should never end up with certainty about some controversial moral view. So, in our decision-analytic language, we should be thinking about imperfect information—information that improves our epistemic state rather than giving us certainty—instead of perfect information.

That we will gain only imperfect information doesn't change the framework, but it does make the mathematics more difficult, and it does reduce the expected utility of gaining new information. To illustrate how the framework works for imperfect information, consider the *Salesman* example again, and suppose that rather than being able to gain perfect information about the weather, Jonny is able to gain only imperfect information by asking a semi-reliable market forecaster: if it will be sunny, the forecaster will say so 90% of the time; if it will rain, the forecaster will say so 90% of the time.

As with perfect information, the expected utility of gaining new information is given by the expected utility of one's decision given that new information minus the expected utility of one's decision without that information. In order to work this out, we first must work out what credences Jonny ought to have, depending on what the forecaster says. Using Bayes' theorem, if the forecaster says it will be sunny, then Jonny ought to believe with 90% certainty that it will be sunny; similarly, if the forecaster says it will rain, then Jonny ought to believe with 90% certainty that it will rain.¹¹

The expected utility of Jonny hearing the forecaster's view is therefore:

$$\begin{aligned} & \text{Expected utility of best decision given new information} - \text{Expected utility} \\ & \text{of best decision given no new information} \\ & = (0.5(0.9 \times (\$50 - 10) + 0.1 \times (-\$10)) + 0.5 \times \$0) - (0.5 \times (\$50 - 10) + 0.5 \times \\ & (-10)) \\ & = \$17.50 - \$15 \\ & = \$2.50 \end{aligned}$$

So Jonny should be willing to pay up to \$2.50 to hear the forecaster's opinion.

That was an illustration of the expected utility of gaining imperfect empirical information. For an example of the expected choice-worthiness of gaining imperfect moral information, consider again the philanthropic foundation example. As before, the foundation's leader has 99% credence that improving conditions in factory farms is of no value, and 1% credence

¹¹ In what follows, we'll only use examples of imperfect information sources where there is an equal probability of Type-I and Type-II errors. Incorporating the idea that some information sources might be more likely to make a Type-I than a Type-II error again makes the math slightly more complicated but would not change the framework for estimating the value of imperfect information, so we leave it out for simplicity. It's an interesting question, beyond the scope of this chapter, whether philosophical argument is more likely to make Type-I or Type-II errors regarding the wrongness of an action.

that improving those conditions for one hen is 1/100th as good as providing one QALY. She does not know the philosophical literature on animal ethics, but could investigate it to decide what answer the weight of philosophical argument favours. Suppose that the foundation leader believes that the majority opinion of the philosophical community will correctly assess moral issues 90% of the time.

If so, then she should estimate that, upon investigating the philosophical literature, there is 10.8% chance¹² that the weight of philosophical argument will favour the idea that improving conditions on factory farms matters morally, and 89.2% chance that it will favour the idea that improving conditions on factory farms does not matter morally. In accordance with Bayes' theorem, if she investigates the philosophical literature and the arguments favour the idea that improving conditions on factory farms matters morally, she should come to have a credence of 8.3% that improving conditions on factory farms matters morally; if she investigates the philosophical literature and the arguments favour the idea that improving conditions on factory farms does not matter morally, she should come to have a credence 0.11% that improving conditions on factory farms matters morally. Providing the grant for cage-free hens only has the higher expected choice-worthiness if she finds that the philosophical arguments favour the idea that improving conditions on factory farms matters morally. So the expected choice-worthiness of gaining this imperfect information is equivalent to providing $10.8\% \times ((8.3\% \times 3,800,000) - 100,000) = 23,263$ QALYs. Because she can provide a QALY for \$100, the foundation leader should therefore be willing to spend \$2.3 million to gain this imperfect information in order to have a better estimate of how to spend the \$10 million. Or, alternatively, 18.5% of her budget, if that budget is fixed (down from 27% for perfect information).

The extent to which the fact that moral information is inevitably imperfect information reduces the expected choice-worthiness of new information depends on how reliable or unreliable we believe the information to be. If the information is fairly reliable—we believe that the weight of philosophical argument is correct 90% of the time—then the expected choice-worthiness of gaining imperfect moral information can still be high. In contrast, if the foundation leader believed philosophical argument to be only 75% reliable,

¹² A $1\% \times 90\%$ chance that the weight of philosophical argument favours the moral importance of improving conditions on factory farms and this is the correct view, and a $99\% \times 10\%$ chance of a false positive.

then the expected choice-worthiness of gaining that imperfect moral information would be only \$298,860, or 2.9% of her budget.¹³ If the foundation leader believed philosophical argument to be only 70% reliable, then, for the purposes of this decision, the expected choice-worthiness of gaining the new moral information would be nil, because no matter which answer the philosophical literature favoured, she would not change her decision to fund the Against Malaria Foundation.¹⁴

The expected choice-worthiness of gaining new moral information depends crucially, therefore, on how reliable one takes the information to be. If one believes that one will not learn very much from doing study, research or reflection on ethical matters, then the expected choice-worthiness of gaining that moral information will be low.

But, at least sometimes, ethical study and reflection can result in drastic changes to one's beliefs, in ways that seem epistemically warranted. Many people, for example, have on the basis of philosophical arguments moved from having high credence that donating a large proportion of one's resources to effective causes is merely supererogatory to having high credence that doing so is obligatory. For these people, the expected choice-worthiness of the imperfect information they gained from engaging with philosophical arguments and personal reflection was not very different from the expected choice-worthiness that perfect information would have had.

It is, of course, difficult to assess the reliability of studying or researching moral philosophy, or engaging in ethical reflection. To get a crude approximation of the expected choice-worthiness of imperfect information, however, one could ask oneself: after a certain time period of investigation,

¹³ In this case:

$$P(\text{Philosophy favours animals}) = 0.75 \times 0.01 + 0.25 \times 0.99 = 0.255$$

So:

$$P(\text{Animals matter} | \text{Philosophy favours animals})$$

$$= \frac{P(\text{Philosophy favours animals} | \text{Animals matter}) \times P(\text{Animals matter})}{P(\text{Philosophy favours animals})}$$

$$= 0.75 \times 0.01 / 0.255 = 0.0294$$

$$\text{Gaining this imperfect information, therefore, is worth } 0.255 \times (0.0294 \times 3,800,000 - 100,000) = 2988.6 \text{ QALYs or } \$298,860.$$

¹⁴ In this case:

$$P(\text{Philosophy favours animals}) = 0.7 \times 0.01 + 0.3 \times 0.99 = 0.304$$

$$\frac{P(\text{Animals matter} | \text{Philosophy favours animals})}{P(\text{Animals matter})} = \frac{P(\text{Philosophy favours animals} | \text{Animals matter})}{P(\text{Philosophy favours animals})} = 0.7 \times 0.01 / 0.304 = 0.023$$

Even in the situation where the foundation leader learns that the weight of philosophical argument favours animals mattering, the expected choice-worthiness of giving the grant to the cage-free campaign is only worth as much as $0.023 \times 3,800,000 = 87,400$ QALYs, which is less than could be gained from giving the grant to the Against Malaria Foundation.

how likely am I to have changed my moral view (in a way that is epistemically warranted)? And, given that I change my view, what is the difference in choice-worthiness between the decision I'd make then and the decision I'd make now? This procedure would approximate the value of information, but it wouldn't be quite satisfactory. Really, you'd want to provide a probability distribution over all the possible ways in which you could change your view, and the gain in value for all of those possibilities. The expected choice-worthiness of imperfect information would be the integral of the gains in choice-worthiness with respect to that probability distribution. This would be very hard to calculate exactly, but for most cases it suffices to point out that it is quite large rather than to calculate it accurately.

How could you even guess the likelihoods of changing one's view? A simple way would be to use induction from past experience: if one has already spent a fair bit of time doing ethical research, one could look at how many months one had spent doing the research, how many times one had changed one's view on the topic, and how big a difference to the expected value of one's decisions those changes made. This would give one some amount of data by which to make a guess about how likely it is for one to change one's view given additional research. And if one hasn't done research in the past, then one could use information about the likelihood of change from those who have.¹⁵

¹⁵ One could construe such belief-changes in an alternative way: that they are evidence of *overconfidence*, rather than rational updates on the part of the decision-maker (we thank Christian Tarsney for this objection).

We agree (as argued in Chapter 1) that people are often morally overconfident, and that they probably often over-update on new moral considerations. This might sometimes attenuate the apparent impact of gaining new moral evidence from studying moral philosophy. However, overconfidence can also lead one *not* to change one's view even though one ought to have done so, and it seems that self-serving, conformity, and status quo biases all make people more resistant to changing their moral beliefs than they ought to be.

What's more, it's clear that one can gain significant moral information through either moral philosophy or lived experience or both. If one studies moral philosophy, then one is exposed to a strictly larger range of arguments, and therefore evidence, than one otherwise would have; if one has a larger and more diverse array of life-experiences than one is, again, exposed to a larger set of evidence than one would otherwise have been.

Further, there are some moral views that seem to be more common among philosophers. The proportion of philosophers who believe it is wrong to eat factory-farmed meat, for example, is much higher than the proportion of the general public who believe the same. The same is true, we believe, for the idea that those who live in rich countries have significant obligations to strangers.

Conclusion

In this chapter we provided a framework for assessing the expected choice-worthiness of gaining new moral information, and illustrated this framework with respect to the decisions facing a philanthropic foundation and a young altruistically minded person deciding what career to pursue. Though conclusions on these matters are necessarily tentative, depending crucially on the credences of the decision-maker, it seems to us that, in at least some situations, the expected choice-worthiness of engaging in further ethical reflection, study, or research can be very high.