

Ethics, Evolution, and the Coincidence Problem: a Skeptical Appraisal

Abstract:

It may seem to require an extraordinary coincidence if the moral beliefs that have evolved under natural selection turn out also reliably to be true, given that matters of truth and falsity are irrelevant in accounting for the selection-pressures shaping human moral psychology. According to Street, this *Coincidence Problem* represents a serious objection to meta-ethical realism. I identify a number of problems facing her argument, suggesting that the appearance of a problem is misleading: a commitment to a coincidence of this kind need not be unacceptable for the realist. Central to my argument will be an analogy between the *Coincidence Problem* and the *Fine-Tuning Problem* in the philosophy of cosmology. I show that the seriousness of the *Coincidence Problem* as an objection to realism depends ultimately on controversial questions in the philosophy of probability, Bayesian epistemology, and the philosophy of religion.

1. Introduction

Evidence suggests that human morality has evolved under natural selection.¹ However, our best understanding of the evolutionary process indicates that there has been no selection for true moral beliefs: if some element of human moral psychology has evolved under selection, the advantages conferred by this trait are explanatorily independent of the accuracy of any corresponding moral judgments.² In light of this, many philosophers feel that evolutionary explanations for human morality are debunking: they provide defeaters for

¹ Bowles & Gintis (2011), de Waal (1996), Haidt & Kesebir (2010), Joyce (2006), Richerson & Boyd (2005), Sinnott-Armstrong (2008), Sober & Wilson (1998).

² Haidt & Kesebir (2010, 808-814); Street (2006, 125-135).

the beliefs whose origins they explain.³ A wide array of arguments have been taken to support this conclusion: some appeal to Ockham's Razor,⁴ others to considerations of Nozickian sensitivity.⁵ My focus in this paper will be on a line of argument centred on what I call the *Coincidence Problem*.

Very roughly, the problem is as follows. It may seem to require an extraordinary coincidence if those moral beliefs that have evolved under natural selection turned out also reliably to be true, given that matters of truth and falsity are irrelevant in accounting for the selection-pressures shaping human moral psychology. Since we cannot reasonably expect such a coincidence, we might think, we ought to suspend judgment with respect to those beliefs.

This *Coincidence Problem* is most prominent in a series of recent articles by Sharon Street,⁶ though Street believes it arises only if we assume some form of meta-ethical realism. Street herself endorses meta-ethical constructivism.⁷ The problem also features in work by other philosophers hoping to derive rather different conclusions,⁸ but I will not provide detailed consideration of their work in this paper. Within the present debate, the *Coincidence Problem* is, I believe, the most promising articulation of the worry that evolutionary explanations threaten to debunk our moral beliefs. Street's framing of the problem is by far the most sophisticated and influential. However, I believe her argument is

³ de Lazari-Radek & Singer (2012), Crisp (2006), Huemer (2008), Joyce (2001, 2006), Kahane (2011), Parfit (2011), Rosenberg (2011), Ruse (1986), Singer (1981, 2005, 2006). Kahane (2011) and Street (2006, 2008a, 2011) hedge this conclusion: evolutionary explanations are debunking iff we assume some form of realism.

⁴ Joyce (2006), Ruse (1986); cf. Gibbard (1990).

⁵ Joyce (2001), Ruse (1986), Parfit (2011). Cf. Nozick (1981, 342-348).

⁶ Street (2006, 2008a, 2011).

⁷ Street (2008b).

⁸ Huemer (2008), Rosenberg (2011).

subject to a number of problems that have so far escaped notice. The suggestion that we cannot reasonably believe that natural selection just happened to favour the evolution of objectively correct ethical beliefs may be intuitively compelling, but we shouldn't rest content with our intuitions on the matter. It is well-established that our intuitions about coincidences are unreliable: we are far too quick to think that the regularities we find can't be merely accidental.⁹ We should consider more carefully whether our intuitions stand up to scrutiny. By drawing an analogy between the *Coincidence Problem* and the *Fine-Tuning Problem* in the philosophy of cosmology, I show that there are many reasons to believe that our intuitions may have led us astray in the case of evolution and ethics.

In the next section, I'll set out Street's argument that the *Coincidence Problem* represents a serious problem for meta-ethical realism. In section 3, I'll examine whether realists are forced to posit some form of coincidence in order to maintain that there has been a correlation between ethical truth and selective advantage, and whether our preference for realism over constructivism really makes any difference on this point. The answers are 'yes' and 'no', respectively. In sections 4 through 7, I consider whether it is really unreasonable to posit a coincidence of this sort. I point to a number of considerations suggesting that this may be unobjectionable for the realist. An analogy between the *Coincidence Problem* and the *Fine-Tuning Problem* will be used to illuminate both strengths and weaknesses of the *Coincidence Problem*. I am inclined to believe that the weaknesses ultimately outweigh the strengths, but establishing that conclusion is beyond the scope of my paper, requiring us to resolve controversial questions in the philosophy of probability, Bayesian epistemology, and the philosophy of religion.

⁹ Kahneman (2011, 114-118), Zusne & Jones (1989).

2. Street's "Darwinian Dilemma"

The aim of this section is to set out Street's argument so as to make clear why she believes the *Coincidence Problem* represents a challenge for realism and why she supposes this challenge can be met if we choose constructivism as our preferred meta-ethical view.

We can begin by noting that, as debunking arguments go, Street's is atypical in two respects. Firstly, most philosophers working in this area are principally interested in debunking arguments targeting our moral beliefs. Street's stated concern is with beliefs about practical reasons in general. To take account of this more expansive target, I'll speak in terms of 'ethical beliefs' or 'ethical facts', taking this to encompass beliefs or facts about practical normativity, including but not limited to morality.

Street believes that our ethical beliefs and intuitions are influenced to a great extent by certain innate biases: so much so that rejecting those of our beliefs which have been influenced in this way would be catastrophic from the perspective of our having any reasonable grip on how we ought to live. The second notable feature of her argument is that she thinks meta-ethical realists - and *only* realists - find difficulty in accepting evolutionary accounts of our ethical beliefs without being forced into this kind of catastrophic normative skepticism.

'Realism' here denotes the view that practical reasons are constitutively independent of our evaluative attitudes.¹⁰ Street's own view is a form of *meta-ethical constructivism*. She believes that normative facts are constructed out of our attitudes: facts about what a person ought to do are constituted by facts about the reasons she would self-ascribe when her normative judgments are brought into reflective equilibrium.¹¹ On this picture, value arises out of our valuing. Street believes anyone accepting a view of this kind faces no difficulty in

¹⁰ This is taken by Street (2011) to include quasi-realist expressivists like Blackburn (1998) and Gibbard (2003).

¹¹ Street (2008b).

accommodating evolutionary explanations of our ethical beliefs. Other philosophers have sought to formulate evolutionary debunking arguments that carry catastrophic implications regardless of whether we suppose the ethical facts are objective or constructed.¹²

With these preliminary points in mind, let's now set out Street's argument. We start off from the assumption that "the forces of natural selection have had a tremendous influence on the content of human evaluative judgements." (Street 2006, 113) Given this assumption, Street confronts the realist with a 'Darwinian dilemma': "either the evolutionary influence tended to push our normative judgments *toward* the independent normative truth, or else it tended to push them *away from* or *in ways that bear no relation to* that truth." (2011, 12) In other words, of the normative beliefs favoured by natural selection, either these have tended to be true, or else they have not.

Given the assumption that our fundamental normative principles have been determined by natural selection, the denial of such a tendency seems clearly problematic: our normative outlook would be saturated with error. This might not be so bad if we had some independent means of weeding out these errors, but Street (2006, 124) denies the existence of any such capacity.

Suppose, then, that the realist takes the second horn: she posits that there *has* been a consistent correlation between those affective dispositions which have been fitness-maximizing and those which are conducive to believing mind-independent normative truths. Then, Street insists, the realist has some explaining to do. "This degree of overlap between the content of evaluative truth and the content of the judgements that natural selection pushed us in the direction of making begs for an explanation." (2006, 125) This explanatory burden, Street suggests, cannot be met: "the realist about normativity owes us an explanation of this striking fact, but has none" (2008a, 207).

¹² See Joyce (2001, 2006, forthcoming).

Street maintains that the only explanatory strategy open to the realist is to adopt what she calls the *tracking account*: the view that there has been selection for psychological dispositions that result in true normative beliefs. Street maintains that, for the realist, “the tracking account is the only (non-coincidence-positing) way of seeing how evolutionary forces could have pushed our values *toward* independent normative truths” (2011, 13). The tracking account is opposed to what she calls the *adaptive link account*, according to which matters of truth and falsity are irrelevant in explaining the fitness-benefits associated with these evaluative dispositions: such adaptive advantages as they conferred are explained entirely by their motivational effects and by the independent fitness-benefits associated with the behaviours they motivate. Street rejects the tracking account: it is, she says, “scientifically indefensible.” (2011, 13)

Street concludes that realists cannot grasp the second horn of the dilemma, having no means to explain the existence of a consistent correlation between those affective dispositions which have been fitness-maximizing and those which have been conducive to believing normative truths. Anti-realists, she argues, face no similar difficulty, since they count the normative facts as a function of our attitudes, and our attitudes are a function of our evolutionary history:

Antirealism explains the overlap not with any scientific hypothesis such as the tracking account, but rather with the metaethical hypothesis that value is something that arises as a function of the evaluative attitudes of valuing creatures – attitudes the content of which happened to be shaped by natural selection. (Street 2006, 154)

3. Coincidences

Is Street correct to suppose that realists who take the second horn in her dilemma are committed to an implausible coincidence involving natural selection? In assessing this

charge, it's worthwhile to begin by considering exactly what a coincidence is supposed to be. With this complete (or as complete as need be), we'll then consider whether a coincidence has to be posited by the realist who accepts the falsity of the tracking account, and whether the anti-realist is able to escape the same commitment. The answers will be 'yes' and 'no', respectively.

The key upshot of my argument in this section is as follows. It has been thought that the *Coincidence Problem* represents a challenge for realism insofar as the realist is unable to explain why selection has reliably favoured true ethical beliefs, the assumption being that explaining this relationship represents a key desideratum for a meta-ethical theory and one the constructivist can easily meet.¹³ If my argument is correct, this conception of the problem is mistaken: either there is no especial problem facing the realist or else the problem is not to do with finding suitable explanatory links between selection and ethical truth. The next section of this paper will explore a better representation of the problem, involving Bayesian probabilities.

3.1 Conceptual analysis

In trying to understand more exactly what it means to describe something as 'a coincidence', I want to focus on the relation between coincidence and probability. It is sometimes proposed that nothing is properly called a coincidence unless it is sufficiently unlikely to have occurred. Thus, Hart and Honoré write:

We speak of a coincidence whenever the conjunction of two or more events in certain spatial or temporal relations is (1) very unlikely by ordinary standards and (2) for some reason significant or important, provided (3) that they occur without human contrivance and (4) are independent of each other. (1959, 74)

¹³ See esp., Enoch (2010, 427-429), Schafer (2010, 483-485), Street (2006).

However, not everyone uses the term in this way. David Owens (1992) proposes that we define ‘coincidence’ purely in terms of Hart and Honoré’s fourth condition: “A coincidence is an event which can be divided into components separately produced by independent causal factors.” (Owens 1992, 13)¹⁴ On this interpretation, coincidences do *not* have to be chancy or surprising.

Owens suggests that Hart and Honoré capture the term in its ordinary meaning and that his represents a reforming definition. Owens’ definition is nonetheless highly faithful to our linguistic intuitions. Suppose I were to say: ‘I pray every afternoon for the sun to rise the next day, and it does rise; and that is not simply a coincidence.’¹⁵ The implication here is that my prayer somehow controls the rising of the sun. Obviously, that’s false, and my statement correspondingly outrageous. By contrast, it seems tangential to our acceptance or rejection of my statement whether it was unlikely or surprising that sunrise would follow my prayers.

In the same vein, we might consider some of the examples noted by Marc Lange (2010) in his discussion of mathematical coincidences. According to Davis (1981), it is just a coincidence that 9 is the 13th digit in the decimal expansion of both π and e . According to Corfield (2004), it is just a coincidence that both 13 and 31 are prime. It is difficult to see in what sense these facts could be counted as improbable or surprising.

Of course, mathematical facts are not (and do not describe) events whose causes may be independent or overlapping. In that sense, Owens’ definition also falters here. However, as Lange (2010, 317) suggests, we can account for the existence of mathematical coincidences by way of a definition modelled on Owens’, which speaks in terms of *facts*

¹⁴As Owens notes, Hart and Honoré’s third criterion is in fact a trivial consequence of the fourth. Thus, his account rejects only (1) and (2).

¹⁵ I base this example on a case discussed by Owens (1992, 8).

rather than *events* and *explanation* rather than *causation*: a conjunction of mathematical facts is counted as coincidental insofar as the conjuncts are suitably explanatorily independent of one another.

Arguably, the concept *coincidence* does not admit of any neat definition by way of necessary and sufficient conditions. Virtually no concept does. A more fruitful approach may be to think of the concept as constituted by a *prototype*: a list of weighted features used to determine category membership.¹⁶ I suspect that Hart and Honoré capture the features summarised under the *coincidence* prototype. However, Owens' reforming definition ultimately provides a better approximation of our ordinary meaning, since the weighting of these features seems to strongly favour explanatory independence as the key deciding factor. Explanatory independence is in this sense the heart of the vernacular *coincidence* concept.

Insofar as we mean to be more exact in our talk, I would suggest that we treat a coincidence as being simply a conjunction of facts whose conjuncts are explanatorily independent of one another: neither fact figures in the explanation of the other, and there is no relevant¹⁷ explanatory factor shared by the members of the conjunction. From here, I'm going to use the term 'coincidence' under this purely explanatory definition. It's not especially important to me that we adopt this convention, however. What *is* important is

¹⁶ Rosch (1973).

¹⁷ Typically, the sorts of things which are described as coincidences will share some remote explanatory factor(s). For example, I share my birthday with Peter Singer, and this, we would say, is just a coincidence. But these events have the Big Bang as a common cause. To make sense of this, we need to attend to the pragmatics of explanation. In certain contexts, factors which contribute to the explanation of a certain phenomenon are nonetheless rightly ignored. For example, if a student is asked whether there are any common causes of the First and Second World War, she would do well to cite German nationalism, but not the origin of humanity in the East African Rift Valley. The conjunction of two facts which share certain explanatory factors may be rightly described as a coincidence if we are properly ignoring those explanatory factors shared by the conjuncts.

that we not get confused. We should take care to keep one mental file for the question of whether there might be some suitable explanatory connection between the (mind-independent) ethical facts and the existence of certain selection pressures, and a separate file for the question of whether such an alignment might be improbable or surprising. We shouldn't presume that the two questions must be answered in the same way.

3.2 Is a coincidence required?

In deciding whether an alignment between accuracy and adaptive advantage would have to involve a coincidence, Street seems to focus to a large extent on the existence of suitable explanatory connections linking the right values to greater relative fitness. Thus, the explanatory conception of coincidence that I've just outlined seems to fit the terms of the debate. Using this conception, let's then ask whether it would really have to be a coincidence if natural selection has reliably favoured the evolution of ethical intuitions and judgments that are also objectively true. Having decided this question, we'll then consider whether rejecting meta-ethical realism in favour of constructivism makes any difference to our need to posit a coincidence of this kind.

One prominent objection that has been raised against Street is that she ignores the possibility of indirect explanatory links.¹⁸ Street appears to consider only two ways in which to connect accuracy and adaptive advantage: either the accuracy of certain ethical intuitions explains why they have been selected for, or else it's the other way round. There is also the possibility that one and the same set of factors might explain *both* why certain ethical norms have proven adaptively advantageous *and* why those norms are correct. Thus, it needn't be a coincidence, even for the realist, that certain norms have been both

¹⁸ See Brosnan (2011), Enoch (2010), Schafer (2010), Skarsaune (2011), Wielenberg (2010).

reproductively advantageous and correspondent to the moral truth: there may be an indirect explanatory connection linking the two.

I don't think it's especially damning that this possibility has been neglected. The proposal succeeds only in relocating the point at which a coincidence must be posited. Suppose that one and the same factor explains why certain ethical norms have proven adaptively advantageous *and* why those norms are correct. Then we should ask: is it just a coincidence that the factors which have made norms fitness-raising are also the very same factors which make norms correct?¹⁹ Suppose, for the sake of argument, that (a) the tendency of certain norms to promote a peaceful, cooperative social life explains why such norms have raised the relative fitness of groups upholding them and (b) the tendency of the same norms to promote a peaceful, cooperative social life makes them correct. Then, we must ask: is the conjunction of (a) and (b) just a coincidence? Presumably, if we are otherwise impressed by Street's argument, we won't be inclined to suppose that (b) explains (a) or vice versa. To avoid postulating a coincidence, we might then posit yet another third factor, which accounts for both (a) and (b). This would just displace the problem one step again. Unless we're happy to countenance an infinite regress of third-factor explanations, the appeal to indirect explanatory connections fails to rule out that a coincidence had to obtain. For this reason, I think Street is correct to insist that "the tracking account is the only (non-coincidence-positing) way of seeing how evolutionary forces could have pushed our values *toward* independent normative truths" (2011, 13).

¹⁹ Enoch (2010, 433) notes this problem. His response involves a switch of mental files: he argues simply that this deeper alignment is not plausibly regarded as having a low objective probability. Probabilistic issues of this kind are addressed in section 5, where I argue that they are insufficient to defuse the *Coincidence Problem*. On the conception of coincidence we are using here, Enoch's probabilistic reflections are simply irrelevant in deciding whether a coincidence had to obtain.

So much for the realist. I'm now going to argue that constructivists do *not* avoid having to say that only by a coincidence has natural selection favoured the evolution of ethical intuitions/beliefs that happen also to be true. I believe that Street's argument to the contrary is flawed.

If there's to be some explanatory connection between accuracy and adaptive advantage, it seems there are only three possibilities: certain beliefs have been selected for because they are true; certain ethical beliefs are true because they were selected for; or else there exists some indirect explanatory connection. The first possibility is ruled out by the adaptive link account and the third doesn't avoid the need to posit a coincidence, as we've just seen.

This leaves only the second option: certain ethical beliefs are true because holding beliefs with such contents raised the relative fitness of our ancestors. If the constructivist is to avoid the commitment to a coincidence, she must take this view. However, this view seems implausible. On its face, the fact that a given ethical belief was selected for does nothing to make it true. For example, it may be that belief in the justice of retributive punishment can be explained in terms of natural selection.²⁰ However, the fact that a disposition to favour retributive punishment helped our ancestors contribute more descendants to subsequent generations seems irrelevant to the truth or falsity of retributivism as a theory of punishment. Anyone who tried to argue that retributivism must be true by appeal to such considerations would find a hard time getting a foothold in the punishment literature.

In suggesting that the truth-value our ethical beliefs is independent of whether holding such beliefs led to increased reproductive success in the Pleistocene, you may worry that I am begging the question against the constructivist. This worry is misplaced. Constructivism itself seems inconsistent with the view that the truth of our evaluative beliefs is grounded in facts about past selection.

²⁰ See Hoffman (2014).

According to constructivism, a belief of the form *p is a reason for S to ϕ* is true iff and because a belief with that content coheres suitably with *S*'s other normative judgments and so would be held in reflective equilibrium. Suppose I believe that I ought to punish those who do wrong, even if nothing good would come of it. On the constructivist theory, it doesn't seem that the truth-value of this belief depends on how it originated. All that matters is how it relates to my other beliefs. If the belief coheres suitably with my other normative judgments, it is true; it doesn't matter whether my holding this belief is due to selection, drift, pleiotropy, phyletic inertia, or some other cause. On the other hand, if the belief does not cohere, it is false; this also holds regardless of how it originated. Assuming constructivism, the truth-value of my belief depends simply on whether it coheres with my other ethical beliefs. Past selection doesn't enter into it.

Why might we think that constructivism allows us to count the truth of our ethical beliefs as grounded in facts about past selection? We may be tempted to reason as follows. According to the constructivist, my belief in some first-personal normative proposition, *p*, is true because I would accept *p* in reflective equilibrium. But the fact that I would accept *p* in reflective equilibrium is explained, one might think, by the fact that I belong to a species in which the disposition to accept *p* has been favoured by natural selection.²¹ Therefore, the fact that a disposition to accept *p* has been favoured by natural selection explains the

²¹ This may be doubted. Some philosophers of biology hold the view that natural selection cannot explain the properties of individuals, only populations. For this view, see Pust (2001), Sober (1984). For objections, see Matthen (1999), Nanay (2005), and Neander (1988). This view strikes me as highly plausible when restricted to genetic evolution, because origins essentialism entails that a person cannot exist without their lineage (Pust 2001). However, the view strikes me as implausible when applied to cultural evolution, where origins essentialism creates no similar issue. Cultural selection arguably plays a prominent role in the evolution of morality. See Richerson & Boyd (2005), Sober & Wilson (1998).

correctness of my normative verdict. This, I take it, encapsulates Street's reasoning on the matter.

Here is the problem with this line of reasoning. The inference described above relies on treating explanation as a transitive relation. We are asked to conclude that selection explains why I am correct to believe p because selection explains why I am disposed to believe p , which explains why I would believe p in reflective equilibrium, which explains why I am correct to believe p . However, explanation is not in fact a transitive relation: p may explain q and q may explain r without it being the case that p explains r . It is easy to derive such cases by relying on well-known counterexamples to the transitivity of causation.²³ For example, suppose a hiker is walking in the mountains; the cliff-face above her collapses, sending a boulder hurtling toward her head; at the last moment, she ducks, narrowly avoiding death.²⁴ The fact that the boulder falls toward her head explains why the hiker ducks. The fact that the hiker ducks explains why she survives. However, the falling boulder clearly does not explain why the hiker survives: she survives in spite of that.

Of itself, the existence of such counterexamples needn't be too problematic for Street. The observation that explanation is not a transitive relation is not in general a good objection to inferences that rely on chaining together explanatory dependencies. Suppose Paul has a sweet tooth and so eats large amounts of strawberry ice cream. As a result, he gains weight. In this case, we infer that Paul's sweet tooth explains his weight-gain. It would be silly to object that the inference is faulty because explanation isn't a transitive relation. We might reasonably suppose that explanation behaves by-and-large as if it were a transitive relation and may be assumed to behave as such unless we have positive reasons for

²³ For a general discussion of such counterexamples, see Paul & Hall (2013, 214-244).

²⁴ This example due to Ned Hall is generally thought to be the most convincing counterexample to the transitivity of causation, but there are many others. Schaffer (2012) notes: "Virtually everyone in the causation literature seems now to accept at least some of these counterexamples as genuine" (126).

believing otherwise.

I grant this. However, I believe we have such reasons when it comes to the explanatory claims relied on in Street's argument. Most obviously, we can reiterate the point that her conclusion seems mistaken. The explanatory claims on which she relies otherwise seem plausible. If we chain these claims together, we can infer that the truth of our present ethical beliefs is explained by past selection. However, that conclusion looks to be wrong, even if we assume constructivism. We may therefore suspect that we are dealing with a case of transitivity failure.

In addition, we can note that this is to be expected in light of an independently plausible diagnosis of why transitivity breaks down in cases like that involving the hiker and the boulder. Michael Strevens (2008) argues that transitivity fails in such cases due to a shift in the explanatory framework. Explanations are typically evaluated with respect to a set of background assumptions. Thus, the spark may explain the fire, but only against the assumption that oxygen is present. The total set of background assumptions constitutes our explanatory framework. Strevens (2008, 209-212) suggests that explanation is transitive only relative to a fixed framework: we cannot join together explanations if the background assumptions are not held constant.²⁵

It is easy to see that the relevant background assumptions are not constant for the explanatory claims relevant to the hiker example. In particular, when we consider the claim that the ducking explains the hiker's survival, the fact that boulder is falling toward her head is assumed in the background. It is only given this assumption that her ducking makes sense as an explanation for her survival. In fact, only against the background of this assumption does it make sense to speak of the hiker as *surviving*. However, the boulder falling toward her head is not a background assumption earlier in the story: it features as

²⁵ Cf. Nanay (2005, 1106-7).

the explanandum relative to the collapse of the cliff-face and as the explanans relative to the hiker's duck. The explanatory framework shifts, so transitivity breaks down.

There is a similar shift in framework across the explanatory claims relied on in the argument designed to establish that natural selection can explain the truth of our ethical beliefs if we are permitted to assume constructivism. Here they are again. We assume that past selection explains why I believe p . This explains why I would accept p at the point of reflective equilibrium. Assuming constructivism, the fact that I would accept p in reflective equilibrium explains why I am correct to believe p . Note, then, that in this final step, my believing p is a background assumption: the fact that I would accept p in reflective equilibrium explains why I am correct to believe p only on the assumption that I already believe p . However, my believing p is clearly not a background assumption earlier on: like the falling boulder, it features by turns as explanandum and explanans. Since the explanatory framework is not held constant, we should expect that transitivity breaks down.

We have seen a number of reasons, then, to suppose that Street is wrong to believe that past selection explains the truth of our present ethical beliefs if we assume constructivism. As I've noted, that view seems implausible on its face. Furthermore, the only argument in its support falls apart once we take account of the intransitivity of the explanatory relation. Constructivists are no better off than realists, it seems, when it comes to their ability to postulate suitable explanatory links between past selection and the truth of the ethical beliefs selected for. True beliefs weren't selected for, past selection doesn't explain their truth, and third-factor explanations are no help, as we've seen. A coincidence must be involved, whether we are realists or not.

We might still feel some resistance to this conclusion. It may seem obvious that the *Coincidence Problem* must be less of an issue – or no issue at all – for the anti-realist. Guy

Kahane (2011) has argued that evolutionary explanations simply cannot challenge the reliability of our ethical beliefs if we assume some form of constructivism. He writes:

anti-objectivist views claim that our ultimate evaluative concerns are the source of values; they are not themselves answerable to any independent evaluative facts. But if there is no attitude independent truth for our attitudes to track, how could it make sense to worry whether these attitudes have their distal origins in a truth-tracking process? (112)

We don't have to resist my conclusions in order to agree with Kahane's point. If we're convinced that the *Coincidence Problem* must be less of a problem for the anti-realist, we should infer that there is some important aspect of the problem that we have yet to consider. There is. As I've said, though both issues are primed by the word 'coincidence', we should take care to keep one mental file for the question of whether there might be some suitable explanatory connection between the ethical facts and the direction of selection, and a separate file for the question of whether their alignment might be improbable or surprising. We should consider the first file closed. This leaves open whether there might be some difference between realism and anti-realism when we turn to the second file. That's what I now propose to do.

4. Probability and Surprise

There's an additional reason why probabilities have to enter the picture. Street supposes that the realist cannot reasonably suppose that natural selection happened to favour beliefs whose contents match mind-independent facts about the right and the good. But there is no general ban on believing in coincidences. You can happily go ahead and believe that 9 is the 13th digit in the decimal expansion of both π and e . No one thinks that this kind of coincidence cannot reasonably be expected, or even that the need to posit a coincidence like

this provides a reason to doubt the theories which specify the values of π and e . So why can't we happily think that, as a matter of coincidence, natural selection favoured the evolution of objectively true evaluative beliefs?

Street occasionally describes the overlap between the ethical judgments favoured by natural selection and those we think are true as *striking*. The word 'striking' does a lot of work here, I think. Not all coincidences are striking, but those that are striking invite deep suspicion. It doesn't seem striking that 9 is the 13th digit in the decimal expansions of both π and e . Nobody is surprised or astonished by that. By contrast, it would be astonishing if two people coincidentally ran into one another every Sunday afternoon for the last year in highly varied locales. If we found that two people had run into one another repeatedly in this way, we'd strongly suspect that they were coordinating their activities in some way. The hypothesis that it was all just a big coincidence would be accepted only as a last resort.²⁶

We can make progress in evaluating Street's argument, then, by analysing what this property of 'strikingness' amounts to. Luckily, there exists a small philosophical literature devoted to this issue.²⁷ Within this literature, the most widely respected account is due to Paul Horwich (1982). Like the other parties to the debate, Horwich operates within a Bayesian framework where probabilities represent the degrees of confidence of an ideally rational agent, rather than objective chances existing in nature.²⁸

I'm now going to set out Horwich's approach to the nature of surprises, and then show how it can be applied to the *Coincidence Problem*. If the claims that should be made in applying Horwich's framework to the *Coincidence Problem* hold up, we'll have a basis on

²⁶ I borrow this example from Field (1998).

²⁷ See Harker (2012), Horwich (1982), Schlesinger (1991).

²⁸ Ultimately, the differences between Horwich, Harker, and Schlesinger do not matter very much to my argument in the remainder of this paper, as I'm going to focus on whether the *Coincidence Problem* satisfies the first of Horwich's conditions for surprise and no one disputes the necessity of this condition.

which to say that the coincidence is of a sort that can't reasonably be expected. In addition, the framework will allow us to recover an advantage for anti-realism over realism, for reasons I'll explain.

The first condition on an event being surprising, according to Horwich, is that its occurrence would be assigned a very low prior probability given one's beliefs about the relevant circumstances. Thus, if I have a coin believed to be fair, I should assign a very low prior probability to obtaining heads 100 times in a row - an outcome that would be very surprising. Letting C be the believed proposition *The coin is fair* and E be *The coin is tossed 100 times, yielding heads every time*, we write this as: $\text{Pr}_{\text{old}}(E|C) \approx 0$. This condition, however, is insufficient. We should assign equally low prior probability to any proposition describing the coin as being tossed one hundred times and yielding any particular sequence of roughly equal heads and tails. No sequence of that kind will surprise us, however.

As a second condition, Horwich proposes that there should be some rival interpretation of the relevant set-up, which is not already assigned a sufficiently low prior probability, and relative to which the surprising outcome is in fact antecedently probable. Thus, let K be the proposition *The coin is heavily biased towards heads*. Supposing that K was not regarded as being too implausible at the outset, the following condition ought then to be satisfied: $\text{Pr}_{\text{old}}(K)\text{Pr}_{\text{old}}(E|K) \gg \text{Pr}_{\text{old}}(C)\text{Pr}_{\text{old}}(E|C)$. From this it follows that $\text{Pr}_{\text{old}}(C|E) \ll \text{Pr}_{\text{old}}(C)$.²⁹ Therefore, on learning E , $\text{Pr}_{\text{new}}(C) \ll \text{Pr}_{\text{old}}(C)$.

Here's how we can apply this model to Street's argument. Street thinks the overlap between the ethical judgments favoured by natural selection and those we think are true would be striking if the ethical facts were constitutively independent of our evaluative attitudes. This implies her acceptance of something like the following:

²⁹ See Horwich (1982, 101-102) for the derivation.

Abbreviations:

R	=	the reliability of our evolved ethical beliefs
AL	=	the adaptive link account
MR	=	meta-ethical realism
MC	=	meta-ethical constructivism
B	=	our background knowledge concerning the evolution of humanity

- (i) $\text{Pr}_{\text{old}}(R | AL \wedge MR \wedge B) \approx 0$;
- (ii) $\text{Pr}_{\text{old}}(R | AL \wedge MC \wedge B) \gg 0$;
- (iii) $\text{Pr}_{\text{old}}(MC)$ is not so very low relative to $\text{Pr}_{\text{old}}(MR)$, such that, given (i) and given (ii):
- (iv) $\text{Pr}_{\text{old}}(R | AL \wedge MC \wedge B) \text{Pr}_{\text{old}}(AL \wedge MC \wedge B) \gg \text{Pr}_{\text{old}}(R | AL \wedge MR \wedge B) \text{Pr}_{\text{old}}(AL \wedge MR \wedge B)$.

Are these claims true? Let's focus on (i) and (ii). Although some philosophers might assign a very lower prior to constructivism, I think it would be unreasonable to deny (iii) and (iv) if these first two premises hold up. Furthermore, by reflecting on these premises we'll be able to gain a sense of the sort of advantage for anti-realism that might be on offer in light of this Bayesian re-interpretation of the *Coincidence Problem*.

Firstly, I believe that (ii) is highly plausible: constructivism should lead us to expect our evolved ethical beliefs to be reliable. After all, constructivism treats value as arising out of our valuing: the evaluative facts are a function of our evaluative attitudes, whatever attitudes we happen to end up with. It seems highly plausible, therefore, that our evaluative beliefs cannot be far off from the truth, regardless of their distal origin. This, I take it, is the point of Kahane's rhetorical question: "if there is no attitude independent truth for our attitudes to track, how could it make sense to worry whether these attitudes have their

distal origins in a truth-tracking process?”³⁰ Thus, $\text{Pr}_{\text{old}}(R \mid AL \wedge MC \wedge B)$ should be high.

Realism counts the truth as independent of our attitudes. It is conceptually possible, given realism, that the ethical facts are otherwise than we take them to be, no matter how coherent our beliefs. As Russ Shafer-Landau notes: “Because realism does not see truth as constituted by even idealized attitudes taken towards some subject, realism allows for the possibility that moral truth may elude our best epistemic efforts.” (2003, 225) In addition to being coherent, a reliable belief-set must lock on to a body of mind-independent normative facts. Plausibly, $(AL \wedge B)$ offers no reason to expect that the beliefs favoured by selection will also correspond to this mind-independent domain of fact. Therefore, $\text{Pr}_{\text{old}}(R \mid AL \wedge MR \wedge B)$ should *not* be high, as is implied by (i). Realism introduces an extra condition on reliability, but the evolutionary facts provide no reason to expect our evolved beliefs to satisfy that condition.

We appear, then, to have recovered some form of advantage for constructivism over realism of the sort that failed to emerge when we considered only the explanatory side of things. However, the advantage I’ve just identified is less than what’s needed: I haven’t given us enough to substantiate premise (i). I argued that $(AL \wedge MR \wedge B)$ provides no reason to expect reliability. Premise (i) is stronger: it says that, conditional on $(AL \wedge MR \wedge B)$, we should assign high confidence to unreliability. Whether this stronger claim can be substantiated is an entirely different matter.

In the next section, I’ll assess what has already been said in the literature on the *Coincidence Problem* about the probability that natural selection would favour the evolution of reliable ethical beliefs. This, we’ll see, falls short of settling the matter. It does, however,

³⁰ Cf. Street: “there is no possibility of being ‘off track’ due to evolutionary influences ... *Whatever* our initial set of [evaluative attitudes] might be, what’s ultimately worth pursuing is in some way a function of those.” (2011, 22 n.41)

offer some valuable lessons about the irrelevance of objective probabilities. I'm going to build on these lessons in section 6, where I outline the best – and, I think, only – possible justification for (i).

5. An appeal to chances?

Why should your prior that a fair coin will yield heads on every one of 100 tosses be so low? Because the objective probability of that outcome - $CH(E|C)$ - is very low, and your credence should mirror this: you should defer to chance. This requirement is encapsulated in Lewis's *Principal Principle* (Lewis 1980):

The Principal Principle:

$$\Pr(p | CH(p)=x) = x$$

We might hope to justify assigning a low prior to $(R | AL \wedge MR \wedge B)$ in the same way: by arguing that $CH(R | AL \wedge MR \wedge B) \approx 0$. However, David Copp (2008) argues forcefully against this proposal. He insists that realists are not forced to think it was unlikely that natural selection would favour the emergence of moral beliefs that reliably approximate the truth.

Copp regards probabilistic issues of this kind as central to the problem posed by Street's evolutionary challenge. As Copp understands it, "The basic challenge ... is to explain what it is about the moral truth such that, if the adaptive link account is correct, it is likely that our moral beliefs tend to approximate the truth." (2008, 198) He proposes to meet this challenge as follows. According to his own 'society-centred' view of morality, a moral proposition, p , is true iff and because the moral code that best allows society to meet

its ‘needs’ includes or implies a corresponding norm.³¹ The ‘needs’ of a society are understood to include its need to reproduce itself, for its members to cooperate peacefully, and for peaceful and cooperative relations with neighbouring societies.

Even if the adaptive link account is true, Copp notes, the moral beliefs that we should expect to emerge from the evolutionary process would overlap considerably with those that are true according to the society-centred view: it’s not a matter of chance that natural selection favoured beliefs of that kind. Furthermore, if the society-centred view is true, there’s nothing chancy about that: presumably, it’s a necessary truth, and necessary truths have objective probability 1.³² Thus, it needn’t be a fluke at all if the selection-pressures which have determined our ethical beliefs have favoured beliefs correspondent to the ethical truth: “if our moral beliefs are true or approximately true, this is not a matter of chance.” (Copp 2008, 202)

Roger White has put forward a similar position. He says:

We might deny that it is so unpredictable that evolution should produce creatures with correct moral beliefs. ... [N]atural selection is likely to favor creatures with a sense of obligation toward their offspring. But necessarily, agents *do* have an obligation toward their own offspring. So it is to be expected that evolution will produce creatures with correct moral beliefs on this and a range of other matters.

(2010, 589)³³

What are we to make of this line of reasoning?

³¹ See Copp (1995, 2007).

³² Cf. Enoch (2010), Wielenberg (2010).

³³ Cf. Schafer (2010).

Street, for one, is unimpressed. In reply to Copp, she objects that his argument is trivially question-begging: the question is begged because “our substantive normative views on how we have reason to live are ... merely being taken for granted as at least roughly correct.” (2008a, 215) I presume she would say the very same about White.

I think Street is misreading Copp, but that her complaint is otherwise on the mark. Copp’s conclusion is framed as a conditional: “if our moral beliefs are true or approximately true, this is not a matter of chance.”³⁴ A conditional statement doesn’t presuppose its antecedent, so the truth of our ethical beliefs isn’t being taken for granted here. The charge of question-begging looks to be in better shape when applied to White, who does seem to presume the correctness of our evolved ethical beliefs and certainly makes no effort to hedge his conclusion. As indicated, I think Street is correct to insist that we can’t simply take the reliability of our evolved ethical beliefs for granted in assessing the seriousness of the *Coincidence Problem*. However, this restriction might itself seem question-begging. Unless we’ve already been offered a defeater for these beliefs, why shouldn’t we be able to rely on them? We can explain why using the Bayesian set-up I’ve described. Premise (i) concerns the prior probability of reliable ethical beliefs, conditional on $(AL \wedge MR \wedge B)$. Even granting that we are ordinarily permitted to assume R , we can’t do so when assessing the prior probability of R conditional on $(AL \wedge MR \wedge B)$. That’s just what it means to be talking about the *prior* probability. Thus, we can’t take the reliability of our evolved ethical beliefs for granted in this context.

Let’s return to Copp’s conditional conclusion, which, as I’ve argued, isn’t vulnerable to this problem. Even granting its truth, I think Copp’s position poses less of a challenge to Street’s argument than he thinks. Copp says that if our evolved ethical beliefs are in fact reliable, there was nothing chancy about the evolution of reliable ethical beliefs. This is

³⁴ Cf. Skarsaune’s ‘iffy’ reply to Street: Skarsaune (2010).

entirely compatible with the claim that $\text{Pr}_{\text{old}}(R|AL \wedge MR \wedge B) \approx 0$. Your prior that some set-up will yield a given result could be very low, though you might be very confident that the set-up had a high chance of yielding that result, assuming it does. Here's an illustration. Imagine I pick a coin at random from an assortment of 1000 coins. I know that only one of these coins is heavily biased towards heads, the rest being fair. My prior that my coin will net me one hundred heads in a row should be low; but, my confidence that the coin had a high chance of netting me one hundred heads in a row, conditional on the assumption that it does, should be high.

What Copp's argument manages to do, then, is add greater pressure to the question of how we could justify assigning a low value to $\text{Pr}_{\text{old}}(R|AL \wedge MR \wedge B)$, without providing any positive reason to deny that the value should be low. Granting his conclusion, we can't claim that $\text{CH}(R|AL \wedge MR \wedge B) \approx 0$ unless we assume the falsity of R . Since that assumption is clearly out of bounds in this context, this entails that (i) can't be established by straightforward application of the *Principal Principle* in combination with the premise $\text{CH}(R|AL \wedge MR \wedge B) \approx 0$. In other words, the appeal to objective probabilities in support of (i) is unpromising. If we want to find a plausible basis for this premise, we have to look elsewhere.

6. The Fine-Tuning Analogy

Let's look to the philosophy of cosmology. The fundamental physical constants governing our universe appear fine-tuned to permit the existence of life. Many feel intuitively that this can't merely be a coincidence. This is the *Fine-Tuning Problem* in cosmology. *Prima facie*, there's an analogy here with the *Coincidence Problem*. We can't just assume that the constants happen to fall into the life-permitting region, and we can't just assume that selection happened to lock on to the mind-independent moral facts. This analogy is far from superficial, as I'll show. It can be exploited to construct a plausible basis for (i) - the only

possible basis, I argue. The analogy with the *Fine-Tuning Problem* leads eventually to a series of doubts about the seriousness of the *Coincidence Problem*, which I discuss in the next section.

6.1 *The Fine-Tuning Problem*

As a first step, I'll outline the *Fine-Tuning Problem*, starting with the relevant empirical details. Our universe is governed by a number of fundamental physical parameters, such as the gravitational constant. The values of these parameters cannot be predicted from any physical theory currently known: they are in that sense arbitrary. The values actually taken by these parameters have the interesting property of appearing fine-tuned to permit the existence of life: they fall into a narrowly circumscribed life-permitting range.

The evidence that the fundamental physical parameters are fine-tuned has been thought to support at least one of two extra-cosmic hypotheses: the *Design Hypothesis*, according to which fine-tuning is explained by the existence of a Cosmic Designer; and the *Multiverse Hypothesis*, according to which there exist very many universes, varying randomly in the values of their fundamental physical parameters, such that the probability of at least one life-permitting universe is very high. The support offered to the disjunction of the *Design Hypothesis* and the *Multiverse Hypothesis* by the fine-tuning evidence is typically formalized as a matter of Bayesian confirmation, backed by probabilistic claims akin to those I put forward in section 5.³⁵ The exact formalization varies from author to author. To emphasize the analogy we're interested in, we're going to write it out as follows:

Abbreviations:

K = our background knowledge in physics and cosmology

³⁵ See Collins (2009), Holder (2004), Leslie (1989), Swinburne (1990)

FT = that only if they fall within a narrow range are the constants life-permitting.
 LP = that the constants governing our universe are life-permitting
 LUC = that ours is a lone, uncreated universe, with no designer.
 MV = *The Multiverse Hypothesis*
 D = *The Design Hypothesis*

- (a) $\text{Pr}_{\text{old}}(LP | K \wedge LUC \wedge FT) \approx 0$;
- (b) $\text{Pr}_{\text{old}}(LP | K \wedge (MV \vee D) \wedge FT) \gg 0$;
- (c) $\text{Pr}_{\text{old}}(K \wedge (MV \vee D) \wedge FT)$ is not so low relative to $\text{Pr}_{\text{old}}(K \wedge LUC \wedge FT)$, such that, given (a) and (b):
- (d) $\text{Pr}_{\text{old}}(LP | K \wedge (MV \vee D) \wedge FT) \text{Pr}_{\text{old}}(K \wedge (MV \vee D) \wedge FT) \gg \text{Pr}_{\text{old}}(LP | K \wedge LUC \wedge FT) \text{Pr}_{\text{old}}(K \wedge LUC \wedge FT)$.

From this it follows that $\text{Pr}_{\text{old}}(K \wedge (MV \vee D) \wedge FT | LP) \gg \text{Pr}_{\text{old}}(K \wedge LUC \wedge FT | LP)$. Thus, the discovery that our universe is fine-tuned for life should significantly decrease our confidence that ours is a lone, uncreated universe.

Let's now examine the justification typically offered in support of (a). We'll start by noting that (a) is never supported by appeal to the *Principal Principle* in combination with the posit of some lottery-style cosmogony. No one thinks (a) is true because we have prior reason to suppose the universe came into being via a process in which the values of the parameters were selected at random. Where forthcoming, arguments for (a) appeal instead to the *Principle of Indifference*.³⁶ The *Principle of Indifference* is a means of generating priors from evidential neutrality: cases in which we have no reason to expect one outcome rather

³⁶ Collins (2009), Smolin (1997), Swinburne (1990). Strictly speaking, Collins appeals to a restricted form of the principle, designed to evade problems like the *Water Into Wine* paradox, discussed in section 7.

than another (without necessarily having any positive reason to assign equal expectation to either). The principle says:

The Principle of Indifference:

Necessarily, for any S, P : If S has no more reason to expect any cell - $p_1, p_2, \dots p_n$ - in the partition of a possibility-space, P , S should assign equal probability to any cell, p_i , in the partition: $\text{Pr}(p_i) = 1/n$.

The application to the *Fine-Tuning Problem* appears straightforward. Our background knowledge, K , provides no reason to expect parameter-values falling within any particular range. Given, FT , the range of life-permitting values is known to be very small relative to the space of conceptually possible parameter values. Therefore, by the *Principle of Indifference*, the prior probability of obtaining life-permitting values should be very low, unless we take on some postulate, such as MV or D , which should lead us to expect those values to emerge.

6.2 The Coincidence Problem and the Fine-Tuning Analogy

I'll now show how something like the justification just offered for (a) can be used, *mutatis mutandis*, to support premise (i). There are three steps.

First of all, we need a large possibility-space, corresponding to the vast range of possible parameter settings. In section 3, I said that, conceptually speaking, the mind-independent ethical facts could be totally different from what we've evolved to believe. We might think there's a whole range of ways in which they could be different. Street certainly holds this view. She writes:

According to the normative realist, there are normative truths that hold independently of all our evaluative attitudes. Moreover, as a purely conceptual

matter, these independent normative truths might be anything. In other words, for all our bare normative concepts tell us, survival might be bad, our children’s lives might be worthless, and the fact that someone has helped us might be a reason to hurt that person in return. (2008a, 208)

If we’re inclined to go along with this reasoning, we seem to have the possibility space we need: there is a wide range of conceptually possible worlds in which the ethical facts are varied quite substantially.

As our second step, we need to identify a narrow reliability-permitting region within this space, analogous to the narrow life-permitting region in the space of conceptually possible parameter-settings. To do this, we just need to invoke the principle that selection for or against certain moral norms is independent of their truth or falsity. This entails that in all of those conceptually possible worlds in which the ethical facts are varied from the way we’ve actually evolved to think they are, the direction of selection won’t vary accordingly: natural selection doesn’t track the truth. Since there are so many ways in which to vary the ethical facts, the upshot is a tightly-circumscribed region of the possibility space in which the beliefs favoured by selection line up with the facts.

Finally, we just need to apply the *Principle of Indifference*, appealing to the fact that $(AL \wedge MR \wedge B)$ provides no reason to expect that the moral facts are one way rather than another. Given the size of the reliability-permitting region relative to the whole, we ought to assign a very low prior to the objective ethical facts coinciding with the belief-contents favoured by selection. Therefore, $\text{Pr}_{\text{old}}(R \mid AL \wedge MR \wedge B) \approx 0$.

There’s much to quibble about in this set-up. I’m rather suspicious of the notion of conceptual possibility, for example. I’m just not sure which moral propositions could be true, conceptually speaking. Consider Street’s claim that “as a purely conceptual matter, these independent normative truths might be anything”. One might think there exist *some*

conceptual constraints on what the normative facts could be. For example, one might think it conceptually impossible that the moral facts involve contradictions. Others might attach conceptual necessity to formal principles like the transitivity of betterness or ‘*ought* implies *can*’. Foot (1958) once claimed that there is no moral sense of ‘should’ relative to which someone could believe, as a basic moral principle, that people should not run around trees left-handed or look at hedgehogs by moonlight.³⁷ This might be thought to imply that it’s conceptually impossible for principles like this to be true.

It’s difficult to decide these questions without a better understanding of what is meant by ‘conceptual truth’ and ‘conceptual possibility’. Here is one natural suggestion: p is a conceptual truth iff it’s analytic that p ; q is conceptually possible iff q is logically consistent with the set of conceptual truths. We need then only to find an acceptable account of analyticity. Depending on the conception chosen, this might place very few restrictions on the space of conceptually possible ethical truths. For example, suppose we interpret analyticity as *Frege-analyticity*: a sentence is analytic iff it is synonymous with a logical truth. We would then regard as conceptual truths neither the transitivity of betterness, ‘*ought* implies *can*’, nor the view that examining hedgehogs in moonlight cannot be wrong as a matter of basic moral principle. Suppose instead that we interpret analyticity as *epistemic analyticity* on which a sentence is analytic only if no one could understand the sentence without thereby being disposed to assent to it.³⁸ Timothy Williamson (2006) argues persuasively that virtually nothing is analytic in this sense, including the laws of logic. On either proposal, Street’s claim that there’s an unlimited range of conceptually possible ethical truths appears very plausible.

³⁷ I am indebted to ## for helpful prodding on these points.

³⁸ Boghossian (1996, 2003) champions this conception of analyticity.

As an additional quibble about the set-up I've sketched, we might worry about the use of measure terminology. For the *Fine-Tuning Problem*, modern cosmology allows us to pick out relatively precise real-valued intervals in which a given parameter is life-permitting; the size of the life-permitting range is thus Lebesgue-measurable. It's not so clear what could be meant by the 'size' of the region of conceptually possible worlds in which natural selection and the moral facts are so aligned that some suitable percentage of the beliefs favoured by selection are true. We might feel that we have some intuitive grasp of this notion, but we may also worry that there is no notion to be grasped, intuitively of otherwise: the idea of a measure over the proposed possibility-space is nonsensical.

I will return to these issues in the next section. As we'll see, they are not merely technical quibbles that we can hope to solve at a later date. How these issues are resolved is crucial for whether we should take the *Coincidence Problem* seriously at all. Is there any reason, however, to believe that the kind of set-up I've sketched drives Street's concern with the *Coincidence Problem*? There is. In a later paper, Street writes:

as a conceptual matter, the independent normative truths could be *anything*. For all we know as a conceptual matter, in other words, what's ultimately worth pursuing could well be hand-clasping, or writing the number 587 over and over again, or counting blades of grass. But if there are innumerable things such that it's conceptually possible they're ultimately worth pursuing, ... then what are the odds that our values will have hit, as a matter of sheer coincidence, on those things which are independently really worth pursuing? That the odds seem low is an understatement. (2011, 114)

This strongly implicates acceptance of the kind of set-up I've sketched.³⁹

³⁹ See also Street (2008a, 208).

More importantly, there is, I think, no alternative. I say this because, on the face of it, $(AL \wedge MR \wedge B)$ *does* put us in a state of evidential neutrality with respect to R : we are given no reason to expect that the ethical facts are as we have evolved to think they are, but we are given no positive reason to think that they are not.⁴⁰ Assuming that $(AL \wedge MR \wedge B)$ is in this way neutral with respect to R , there is no way to show that $\text{Pr}_{\text{old}}(R | AL \wedge MR \wedge B) \approx 0$ except by appeal to some means of generating priors from evidential neutrality. This requires reliance on the *Principle of Indifference* in combination with a roomy possibility-space in which the R -region is dwarfed by its complement.

7. ... and the Lord taketh away

In tracing connections between the *Fine-Tuning Problem* and the *Coincidence Problem*, we've so far seen how the latter could be helped along by an analogy with the former. We'll now see how their kinship threatens to undermine the *Coincidence Problem*.

7.1 Problems with the Principle of Indifference

As a first point, we should note that the *Principle of Indifference* is widely rejected.⁴¹ While counterexamples to the principle are routine in introductory courses on probability, I've found surprisingly many moral philosophers unaware that the principle is routinely assigned to the scrapheap of intellectual history.

⁴⁰ White shares this intuition, writing: "evolutionary considerations *fail to vindicate* our moral judgments, they don't provide any *further reason to doubt* the reliability of these judgments." (2010, 590)

⁴¹ See, e.g., van Fraassen (1989): "It is true that the historical controversy extended into our century, but I regard it as clearly settled now that probability is not uniquely assignable on the basis of a Principle of Indifference". (292) Similarly, White (2009) notes: "These days [the *Principle of Indifference*] is commonly dismissed as an old-fashioned item of confusion." (161)

The principle is so widely rejected because there are cases in which it requires us to assign incompatible probabilities to the same outcome under two equivalent descriptions. Richard Von Mises's (1957) *Water into Wine* example illustrates the problem. Let x be the ratio of water to wine in a glass and assume we know only that the value of x lies in the interval $1/2$ to 2 . Then, let x' be the ratio of wine to water in the glass and assume we know only that the value of x' lies within the same interval. The ranges $1/2$ to 1 , 1 to $3/2$, and $3/2$ to 2 are of equivalent size, and so, by the *Principle of Indifference*, should be assigned an equal probability of $1/3$ when considering the value of x . The same reasoning applies to obtaining a value for x' in the ranges $1/2$ to 1 , 1 to $3/2$, and $3/2$ to 2 : by the *Principle of Indifference*, each should be assigned probability $1/3$. However, since x' trivially equals $1/x$, this leads to probabilistic incoherence.

Many philosophers have been led to reject the *Principle of Indifference* in light of such problems, but not all: new solutions to the *Water into Wine* paradox (and other cases like it) are regularly put forward.⁴³ Let's assume, for the sake of argument, that the *Principle of Indifference* is defensible. This only gets us so far. There are other probabilistic issues facing the *Fine-Tuning Problem*. We'll see that these apply, *mutatis mutandis*, to the *Coincidence Problem*.

7.2 Problems with infinite ranges

Much recent work on the *Fine-Tuning Problem* has focused on the *Normalization Problem* and closely associated *Coarse-Tuning Problem*.⁴⁴ The problems arise because it appears that many or all of the parameters governing the universe are, as a matter of conceptual possibility, unbounded (in at least one direction). For example, as a matter of conceptual possibility, the

⁴³ See, e.g., Mikkelsen (2004), White (2009).

⁴⁴ Colyvan et al. (2005), Davies (1992), McGrew et al. (2001), Manson (2000).

gravitational constant could be any real number. The space of conceptually possible parameter-settings is infinite, whereas the life-permitting region represents a finite proportion of the total.

Here is why this creates problems. The *Principle of Indifference* requires us to assign a uniform probability distribution over equal regions of the total possibility space. The total space of conceptually possible parameter-values consists of infinitely many sub-regions of equal size to the life-permitting range. Given the standard axiomatisation of the probability calculus due to Kolmogorov (1950), it is impossible to assign a uniform probability distribution over a countably infinite partition of the total possibility-space consistent with the assumption that probability is a *normalized measure* in which the measure of the total is 1. This is the *Normalization Problem*. It relies on the axiom of *Countable Additivity*:

Countable Additivity:

If p_1, p_2, p_3, \dots is a countable sequence of mutually exclusive outcomes and $\bigvee p_i$ is true iff at least some member of the sequence, p_i , is true, $\Pr(\bigvee p_i) = \sum \Pr(p_i)$.

Though almost always used in mathematical treatments of probability, *Countable Additivity* is one of the more philosophically controversial of the Kolmogorov axioms. It has been thought undesirable because it implies, for reasons just outlined, that there can't be a fair lottery over an infinite set of possible outcomes.⁴⁵ It has been suggested, therefore, that we adopt only the weaker axiom of *Additivity*, which treats the probability of the disjunction of any pair of mutually exclusive outcomes as equal to the sum of their probabilities:

Additivity:

⁴⁵ See de Finetti (1974).

If p and q are mutually exclusive, $\Pr(p \vee q) = \Pr(p) + \Pr(q)$.

This allows for the possibility of fair infinite lotteries. The rejection of *Countable Additivity* for *Additivity* has been endorsed by defenders of the *Fine-Tuning Problem*.⁴⁶ It can be shown, however, that credences which violate *Countable Additivity* leave one vulnerable to a Dutch Book,⁴⁷ typically regarded as the paradigm symptom of irrational credence.⁴⁸

Even if we ignore this issue, we're still in trouble. If probabilities are real-valued, there is only one possible uniform probability distribution over an infinite partition that can be made, consistent with *Additivity* and *Normality*: zero. Any positive value would eventually sum to more than 1. Thus, relative to the unbounded space of conceptually possible parameter-settings, the probability of *any* finite region of the space is always the same no matter how great the region: zero. This creates the *Coarse-Tuning Problem*: the fact that the parameters must be fine-tuned to permit life is no better reason to assign a low prior to the lone, uncreated universe being life-permitting than the fact that there is *some* finite life-permitting region of arbitrary size. The probabilities are the same in either case. It follows that we should be no more surprised that the actual parameter-settings fall within a narrow life-permitting range than that they fall within *any* finite life-permitting range. The existence of *some* upper (and lower) bound on the space of life-permitting universes offers equally good reason to believe in the existence of a God or a multiverse. Many philosophers regard these implications as absurd.

Let me now spell out how the *Coincidence Problem* appears vulnerable to similar objections. The problem arises from the very same source: the apparent lack of any bound

⁴⁶ E.g., Collins (2009).

⁴⁷ Williamson (1999).

⁴⁸ See Howson (2008) for critical discussion of the Dutch Book argument for *Countable Additivity*.

on the possibility space used in arguing for (i). Street claims, as we've seen, that "as a purely conceptual matter, these independent normative truths might be anything." (2008a, 208). Elsewhere she says: "as a conceptual matter ... what's ultimately worth pursuing could well be hand-clasping, or writing the number 587 over and over again, or counting blades of grass. ... [T]here are innumerable things such that it's conceptually possible they're ultimately worth pursuing" (2011, 14) We've seen some plausible interpretations of the notion of conceptual possibility that support this view. In that case, the difficulties we've identified for the *Fine-Tuning Problem* would appear to apply with equal force. In order to attribute any probability to the finite reliability-permitting region, we'd have to reject *Countable Additivity*. We'd have to attribute zero probability to obtaining the reliability-permitting region, and we'd then be faced with our own version of the *Coarse Tuning Problem*. If someone regarded natural selection as being coincidentally accurate in selecting objectively true moral beliefs/intuitions at a rate of 0.1%, we wouldn't think this a commitment to a striking, implausible correlation. By contrast, we might think a reliability-rate of 99.9% too good to be true. However, both the region in which 0.1% of beliefs which are selectively advantageous are true and the region in which 99.9% are true presumably represent bounded regions within the total, unbounded possibility-space, and so end up with the same probability: zero. It would be no more implausible, then, to suppose that natural selection had coincidentally been accurate in 0.1% of cases than in 99.9% of cases. This implication appears equally hard to accept.

Admittedly, it is difficult to say with especial certainty that these problems apply to the *Coincidence Problem* because, as I've noted, it's unclear what is meant by conceptual possibility and similarly unclear how we are to understand the idea of a measure over the relevant possibility space. However, it is no more clear that the argument avoids these issues. Absent some attempt to precisify the notions of conceptual possibility and measure

in such a way that the problems associated with infinite ranges are shown to evaporate, we might reasonably feel unconvinced that (i) is true.

However, someone might think that these problems are less than fatal even if they apply to the *Coincidence Problem*. The *Normalization Problem* and *Coarse-Tuning Problem* have hardly created a consensus in the philosophy of cosmology that fine-tuning is unremarkable. Some are happy to bite the bullets. Jeffrey Koperski (2005) suggests that the *Coarse-Tuning Problem* is just another of the counterintuitive results that arise when we regiment our thinking about infinity. Robin Collins (2009) argues that the existence of an infinite range cannot plausibly undermine the *Fine-Tuning Problem*: this would imply that although the problem is serious given a finite range and more serious the greater the size of that range, the problem disappears altogether when the range becomes infinite. It is more plausible, Collins suggests, to think it should be greater still. Someone could well take a similar bullet-biting stance when it comes to the *Coincidence Problem*.

7.3 Problems with metaphysical non-naturalism

For those so tempted, here is one final respect in which they find a false friend in the *Fine-Tuning Problem*. It is arguable that if evidence of fine-tuning should lead us to revise our cosmological beliefs, then the *Coincidence Problem* shouldn't worry us at all. On the other hand, if the *Fine-Tuning Problem* is really illusory, it's plausible the same holds for the *Coincidence Problem*. The *Coincidence Problem* would be without merit in either case. Let me now explain why we might find these conclusions reasonable.

The problem lies in the status of the *Fine-Tuning Problem* as a putative source of confirmation for theistic religious belief, given that the theistic God may be thought the most natural candidate to fill the role of Cosmic Designer. If theism is probable given our background knowledge, then presumably *R* is to be expected, regardless of *MR* and *AL*. Then (i) would be false. It might be surprising that God chose to impart the correct ethical

beliefs into our species via the process described under the adaptive link account, but that's about it. (He works in mysterious ways.) To the extent that the *Fine-Tuning Problem* provides confirmation for theism, it undermines the *Coincidence Problem*.

Those who suppose that the *Coincidence Problem* is a problem for realism had better, then, have some means of blocking the *Fine-Tuning Problem* as a basis for confidence in theism. However, they're constrained in this respect: the *Coincidence Problem* and the *Fine-Tuning Problem* are sufficiently similar in their fundamentals that many problems facing the latter are likely to apply with equal force against the former. Confirming this prediction, I've pointed to a number of well-known difficulties which face the *Fine-Tuning Problem* and shown that similar problems apply to the *Coincidence Problem*. Either we take problems of this kind seriously, or we don't. If we don't, we're going to need to say something else in response to the *Fine-Tuning Problem*.

This might not seem an especially onerous demand. For example, it may seem obvious that we can just rely on the *Multiverse Hypothesis* to block any inference from fine-tuning to theism. Assuming (a)-(d), *LP* provides strong evidence for $(MV \vee D)$, but this leaves open how we ought to distribute our added confidence across the disjuncts, and whether we ought to end up more confident of *MV* or *D*. Many have felt that we should strongly favour the multiverse over theism.

Against this view, a well-known argument due to Roger White (2000) suggests that *LP* in fact provides no evidence for *MV*. Assuming that the parameter-values obtaining in any member of the multiverse are independent of those obtaining in any other, the probability of obtaining values in the life-permitting range in our universe are no higher conditional on the existence of any greater number of universes. Of course, *MV* raises the probability of life-permitting values in *some* universe. However, rational updating requires conditionalizing on the strongest proposition learned, and the existence of *some* universe

with life-permitting values is weaker than the actual fine-tuning evidence: that the parameters in *this* universe fall within the narrow life-permitting range.

If this sounds unconvincing in the abstract, an example will help. Suppose you are put to a firing squad. Twenty marksmen line up. Each is a fantastic shot: if they mean to kill, there's little chance even one will miss. The guns go off. Astonishingly, you are still alive. You might suspect that the marksmen all just happened to miss. You should certainly give substantial credence to the view that your survival was no accident: the marksmen had decided not to aim for you, or loaded their rifles with blanks. You would not, however, raise your confidence that this firing squad is but one of a collection of similar firing squads, large enough so that at least one squad is expected to miss their target. Although it raises the probability that someone would survive a firing-squad like the one you've faced, the multi-squad hypothesis receives no confirmation from the fact of your survival, whereas the hypothesis of 'design' is rendered significantly more probable. By analogy, *LP* confirms *D*, but not *MV*.⁵⁰

Suppose, nonetheless, that White's argument can be shown to fail and that in assessing the *Fine-Tuning Problem* we are to consider the probability that *some* universe is life-permitting, rather than the probability of *our* universe being life-permitting.⁵¹ Suppose,

⁵⁰ Examples of this kind also help to challenge the suggestion that evidence of fine-tuning should not lead us to revise our cosmological beliefs due to an *observation selection effect* arising from the fact we would not have been able to observe anything had the constants not been life-permitting. It is equally true that you could not have observed anything had you been killed at the firing squad, but this does not undercut the fact that our survival provides evidence that the marksmen aimed to miss (Leslie 1989). Sober (2004) tries to argue against the use of this analogy, but his argument fails for reasons noted by Weisberg (2005). Sober (2009) attempts a reply to these objections.

⁵¹ See Bostrom (2002) and Manson & Thrush (2003) for objections to White.

furthermore, that we therefore assign high confidence to the *Multiverse Hypothesis*. These commitments together raise a different problem.

Notice what happens if we adopt a similar conception of the *Coincidence Problem*. The hypothesis whose probability we're to consider is not that *our* evolved ethical beliefs are reliable but merely that *some* evolved ethical beliefs happen to be reliable. It is not clear why the latter should be thought improbable if we are confident in the *Multiverse Hypothesis*. What is the likelihood that someone somewhere in the many universes happens to have evolved a reliable moral outlook? I don't feel confident that it ought to be very low. Of course, it may still be thought highly unlikely that *our* evolved moral beliefs would turn out to be reliable. It is equally true, however, that the probability that *our* universe would be life-permitting remains as low on the *Multiverse Hypothesis* as on the assumption that ours is a lone, uncreated universe.

In light of these points, we appear to be able to formulate another dilemma. If we think that the right way to frame the *Fine-Tuning Problem* is in terms of the prior probability that *our* universe is fine-tuned for life, the *Multiverse Hypothesis* receives no support from the fine-tuning evidence. By contrast, if this is not the right way to frame the *Fine-Tuning Problem*, it becomes unclear that the prior probability of obtaining reliable evolved ethical beliefs counts as improbable when we consider the *Coincidence Problem*.

There may be some way between the horns of this dilemma. This issue deserves further consideration than I am able to give it here. There are other routes by which to contest the use of fine-tuning as an argument for theism for which I can similarly offer only a cursory treatment. For example, we might insist that even if fine-tuning provides evidence for a Cosmic Designer, it offers little support for the hypothesis of a *theistic* designer: there are many others kinds of designer that we can imagine. Hume (1993 [1779]) famously offers this reply to the old-fashioned *Teleological Argument*. Whether this line of objection is cogent hinges on the prior probability of theism *vis-à-vis* alternative

design-hypotheses. Our prior for theism conditional on the existence of a Cosmic Designer might be significantly higher, in which case the objection would be void.⁵² Another objection that might be lodged against inferring theism from fine-tuning also turns on priors. One might suggest that the absolute prior probability of theism is very low, so its posterior probability remains very low in spite of any evidence offered by fine-tuning.

It is obviously very difficult to say what prior confidence should be assigned to theism without a systematic overview of the philosophy of religion. I cannot offer that here. I can note the following. Nearly 15% of professional philosophers accept or lean toward theism.⁵³ I presume that most would do so regardless of the *Fine-Tuning Problem*. Let's suppose, on a very conservative estimate, that just 10% of professional philosophers would do so. Virtually nobody leans towards any alternative hypotheses involving Cosmic Design; very few accept some form of Deism.⁵⁴ In light of these facts about the views of our peers, one might think that my prior for theism shouldn't be all that low and should be far higher than for any alternative hypothesis consistent with the posit of a Cosmic Designer.

If these admittedly cursory remarks hold up, our options for dismissing the *Fine-Tuning Problem* as confirming theism are severely limited unless we lean on probabilistic objections like those I have noted earlier in this section. However, these objections appear to apply equally well to the *Coincidence Problem*. The *Coincidence Problem* is serious only if the *Fine-Tuning Problem* fails to confirm theism, but if the *Fine-Tuning Problem* fails to confirm theism, this is likely to be because of issues that apply with equal force against the

⁵² In Hume's defence, he offers this objection against the attempt to rest theistic religious belief entirely on evidence of design in nature.

⁵³ Bourget & Chalmers (2010).

⁵⁴ When asked to choose between theism and atheism, less than 1% of respondents surveyed by Bourget & Chalmers chose 'Accept another alternative'.

Coincidence Problem. This is yet another reason not to regard the *Coincidence Problem* as a ground for revising our meta-ethical beliefs in the direction of anti-realism.

Whether it is a decisive reason is a further matter. As I've made clear, the merits of each of the objections noted in this section turns ultimately on deep philosophical questions deserving further consideration than I can provide in this paper. Whether the *Coincidence Problem* ought ultimately to be dismissed will turn on how we decide these important questions.

8. Summary and Conclusion

I have identified a number of considerations that suggest the *Coincidence Problem* may ultimately not be a problem for realism: although realists are committed to a coincidence involving evolution and ethics, this coincidence may not be striking in the way Street supposes.

After having set out Street's argument, I began with an analysis of the *coincidence* concept. I argued that we ought to adopt a purely explanatory analysis and I argued that realists and constructivists are equally committed to the view that only by a coincidence could natural selection have favoured the evolution of ethical beliefs which happen also to be true.

I then sought to understand the probabilistic aspects of the *Coincidence Problem*. Most importantly, I argued that the *Coincidence Problem* needs to be understood as analogous in important respects to the *Fine-Tuning Problem* in the philosophy of cosmology. This analogy has served to indicate a number of problems with the *Coincidence Problem* that have so far escaped notice.

Whether the *Coincidence Problem* should be dismissed in light of these problems is beyond the scope of my argument. Each of the problematic issues I've identified here deserves further consideration and is already associated with a large and growing literature

of its own. What I do take myself to have shown is that the seriousness of the *Coincidence Problem* hinges ultimately on difficult and controversial issues in the philosophy of probability, Bayesian epistemology, and the philosophy of religion.

References:

- Blackburn, S.: 1998, *Ruling Passions*. Oxford: Oxford University Press.
- Boghossian, P.: 1996, 'Analyticity Reconsidered', *Noûs* 30(3), 360-391.
- Boghossian, P.: 2003, 'Epistemic Analyticity: A Defense', *Grazer Philosophische Studien* 66(1), 15-35.
- Bostrom, N. 2002, *Anthropic Bias: Observation Selection Effects in Science and Philosophy*. London: Routledge.
- Bowles, S., and H. Gintis: 2011, *A Cooperative Species: Human Reciprocity and its Evolution*. Princeton, NJ: Princeton University Press.
- Brosnan, K.: 2011, 'Do the Evolutionary Origins of Morality Undermine Moral Knowledge?' *Biology and Philosophy* 26(1), 51-64.
- Collins, R.: 2009, 'The Teleological Argument: An Exploration of the Fine-Tuning of the Universe', in W. L. Craig and J. P. Moreland (eds.), *The Blackwell Companion to Natural Theology* (Oxford: Oxford University Press), pp. 202-281, Blackwell.
- Colyvan, M., J. L. Garfield, and G. Priest: 2005, 'Problems with the Argument from Fine Tuning', *Synthese* 145(3), 325-338.
- Copp, D.: 1995, *Morality, Normativity, and Society*. New York: Oxford University Press.
- Copp, D.: 2007, *Morality in a Natural World*. Cambridge: Cambridge University Press.
- Copp, D.: 2008. 'Darwinian Skepticism About Moral Realism'. *Philosophical Issues* 18(1), 186-206.
- Corfield, D.: 2004, 'Mathematical Kinds, or Being Kind to Mathematics', *Philosophica* 74, 37-62.
- Crisp, R.: 2006, *Reasons and the Good*. Oxford: Oxford University Press.
- Davies, P.: 1992, *The Mind of God: The Scientific Basis for a Rational World*. New York: Touchstone.
- Davis, P. J.: 1981, 'Are there coincidences in mathematics?', *The American Mathematical Monthly* 88, 311-320.
- de Finetti, B.: 1974, *Theory of Probability, Volume 1*, transl. Einaudi. London: Wiley.

- de Lazari-Radek, K., and P. Singer: 2012, 'The Objectivity of Ethics and the Unity of Practical Reason.' *Ethics* 123(1), 9-31.
- de Waal, F. B. M.: 1996, *Good Natured: The Origins of Right and Wrong in Humans and Other Animals*. Cambridge, MA: Harvard University Press.
- Enoch, D.: 2010, 'The Epistemological Challenge to Metanormative Realism: How Best to Understand It, and How to Cope with It', *Philosophical Studies* 148(3), 413-438.
- Foot, P.: 1958, 'Moral Arguments', *Mind* 67(268), 502-513.
- Gibbard, A.: 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.
- Gibbard, A.: 2003, *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- Haidt, J. and S. Kesebir: 2010, 'Morality', In S.T. Fiske, D. T. Gilbert, and G. Lindzey (eds.) *Handbook of Social Psychology, 5th ed.* (Hoboken, NJ: Wiley), pp. 797-832.
- Harker, D.: 2012, 'A surprise for Horwich (and some advocates for the fine-tuning argument (which does not include Horwich (as far as I know)))', *Philosophical Studies* 161(2), 247-261.
- Hart, H. L. A. and A. Honoré: 1959, *Causation in the Law*. Oxford: Oxford University Press.
- Hoffman, M. B.: 2014, *The Punisher's Brain: The Evolution of Judge and Jury*. Cambridge: Cambridge University Press.
- Holder, R. D.: 2004, *God, the Multiverse, and Everything: Modern Cosmology and the Argument from Design*. Aldershot: Ashgate.
- Horwich, P.: 1982, *Probability and Evidence*. Cambridge: Cambridge University Press.
- Howson, C. 2008, 'De Finetti, Countable Additivity, Consistency and Coherence', *British Journal for the Philosophy of Science* 59(1), 1-23.
- Hume, D.: 1993 [1779], *Dialogues Concerning Natural Religion*. Oxford: Oxford University Press.
- Huemer, M.: 2005, *Ethical Intuitionism*. Basingstoke: Palgrave Macmillan.
- Huemer, M.: 2008, 'Revisionary Intuitionism', *Social Philosophy and Policy* 25(1), 368-392.
- Joyce, R.: 2001, *The Myth of Morality*. Cambridge: Cambridge University Press.
- Joyce, R.: 2006, *The Evolution of Morality*. Cambridge, Massachusetts: MIT Press.
- Kahane, G.: 2011, 'Evolutionary Debunking Arguments', *Noûs* 45(1), 103-125.

- Kahneman, D.: 2011, *Thinking, Fast and Slow*. London: Penguin.
- Kolmogorov, A. N.: 1950. *Foundations of the Theory of Probability*, transl. Morrison. New York, NY: Chelsea Publishing Company.
- Koperski, J.: 2005, 'Should We Care About Fine-Tuning?', *British Journal of the Philosophy of Science* 56(2), 303-319.
- Lange, M.: 2010, 'What Are Mathematical Coincidences (and Why Does It Matter)?' *Mind* 119(474), 307-340.
- Leslie, J.: 1989, *Universes*. London: Routledge.
- Lewis, D.: 1980, 'A Subjectivist's Guide to Objective Chance', in R. Jeffrey (ed.) *Studies in Inductive Logic and Probability Volume II*. (Berkeley, CA: University of California Press), pp. 263-294.
- Manson, N. A.: 2000, 'There is No Adequate Definition of 'Fine-Tuned for Life'', *Inquiry* 43 (3), 341-351.
- Manson, N. A., and M. J. Thrush: 2003, 'Fine-Tuning, Multiple Universes, and the 'This Universe' Objection', *Pacific Philosophical Quarterly* 84 (1), 67-83.
- Matthen, M.: 1999, 'Evolution, Wisconsin Style: Selection and the Explanation of Individual Traits.' *British Journal of the Philosophy of Science* 50 (1), 143-150.
- McGrew, T., L. McGrew, and E. Vestrup: 2001, 'Probabilities and the Fine-Tuning Argument: A Sceptical View', *Mind* 110 (440), 1027-1038.
- Mikkelsen, J. M.: 2004, 'Dissolving the Water/Wine Paradox', *British Journal of the Philosophy of Science* 55 (1), 137-145.
- Nanay, B.: 2005, 'Can Cumulative Selection Explain Adaptation?', *Philosophy of Science* 72(5), 1099-1112.
- Neander, K.: 1988, 'What Does Natural Selection Explain? Correction to Sober', *Philosophy of Science* 55(3), 422-426.
- Nozick, R.: 1981, *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Owens, D. J.: 1992, *Causes and Coincidences*. Cambridge: Cambridge University Press.
- Parfit, D.: 2011, *On What Matters, Volume 2*. Oxford: Oxford University Press.
- Paul, L.A. and Hall, N.: 2013, *Causation: A User's Guide*. Oxford: Oxford University Press.

- Pust, J.: 2001, 'Natural Selection Explanation and Origins Essentialism', *Canadian Journal of Philosophy* 31(2), 201-220.
- Richerson, P. J., and R. Boyd: 2005, *Not by Genes Alone: How Culture Transformed Human Evolution*. Chicago, IL: University of Chicago Press.
- Rosenberg, A.: 2011, *The Atheist's Guide to Reality: Enjoying Life Without Illusions*. New York, NY: W. Norton & Co.
- Rosch, E.: 1973. 'Natural categories', *Cognitive Psychology* 4(3), 328-350.
- Ruse, M.: 1986, *Taking Darwin Seriously: A Naturalistic Approach to Philosophy*. Oxford: Blackwell.
- Schafer, K.: 2010, 'Evolution and Normative Skepticism', *Australasian Journal of Philosophy* 88(3), 471-488.
- Schaffer, J.: 2012, 'Grounding, Transitivity, and Contrastivity' in F. Correia and B. Schnieder (eds.) *Metaphysical Grounding: Understanding the Structure of Reality*. (Cambridge: Cambridge University Press), pp. 122-138.
- Schlesinger, G. N.: 1991, *The Sweep of Probability*. Notre Dame, IN: University of Notre Dame Press.
- Shafer-Landau, R.: 2003, *Moral Realism: A Defence*. Oxford: Oxford University Press.
- Sinnott-Armstrong, W. (ed.): 2008, *Moral Psychology, Volume 1: The Evolution of Morality*. Cambridge, Massachusetts: MIT Press.
- Singer, P.: 1981, *The Expanding Circle: Ethics and Sociobiology*. Oxford: Clarendon Press.
- Singer, P.: 2005, 'Ethics and Intuitions', *Journal of Ethics* 9 (3-4), 331 - 352.
- Singer, P.: 2006, 'Review: Richard Joyce, *The Evolution of Morality*', *Notre Dame Philosophical Reviews*: <http://ndpr.nd.edu/news/25012/?id=6383>.
- Skarsaune, K. O.: 2011, 'Darwin and Moral Realism: Survival of the Fittest', *Philosophical Studies* 152(2), 229-243.
- Smolin, L.: 1997, *The Life of the Cosmos*. London: Weidenfeld and Nicolson.
- Sober, E.: 1984, *The Nature of Selection: Evolutionary Theory in Philosophical Focus*. Cambridge, Massachusetts: MIT Press.

- Sober, E.: 2004, 'The Design Argument' in W. E. Mann (ed.) *The Blackwell Guide to the Philosophy of Religion* (Oxford: Blackwell), pp. 117-147.
- Sober, E.: 2009, 'Absence of Evidence and Evidence of Absence: Evidential Transitivity in Connection With Fossils, Fishing, Fine-Tuning, and Firing Squads', *Philosophical Studies* 143(1), 63-90.
- Sober, E., and D. S. Wilson: 1998, *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge, MA: Harvard University Press.
- Street, S.: 2006, 'A Darwinian Dilemma for Realist Theories of Value', *Philosophical Studies* 127(1), 109-166.
- Street, S.: 2008a, 'Reply to Copp: Naturalism, Normativity, and the Varieties of Realism Worth Worrying About', *Philosophical Issues* 18(1), 207-228.
- Street, S.: 2008b, 'Constructivism about Reasons' in R. Shafer-Landau (ed.) *Oxford Studies in Metaethics, Volume 3* (Oxford: Oxford University Press), pp. 207-246.
- Street, S.: 2011, 'Mind-Independence Without Mystery: Why Quasi-Realists Can't Have It Both Ways', in R. Shafer-Landau (ed.) *Oxford Studies in Metaethics, Volume 6* (Oxford: Oxford University Press), pp. 1-32.
- Strevens, M.: 2008, *Depth: An Account of Scientific Explanation*. Cambridge, MA: Harvard University Press.
- Swinburne, R.: 1990, 'Argument from the Fine-Tuning of the Universe' in J. Leslie (ed.) *Physical Cosmology and Philosophy* (London: Macmillan), pp.154-173.
- van Fraassen, B.: 1989, *Laws and Symmetry*. Oxford: Oxford University Press.
- von Mises, R.: 1957, *Probability, Statistics and Truth*, transl. Geiringer. London: George Allen & Unwin.
- Weisberg, J.: 2005, 'Firing Squads and Fine-Tuning: Sober on the Design Argument', *British Journal of the Philosophy of Science* 56(4), 809-821.
- White, R.: 2000, 'Fine-Tuning and Multiple Universes', *Noûs* 34(2), 260-276.
- White, R.: 2009, 'Evidential Symmetry and Mushy Credence' in T. S. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology, Volume 3* (Oxford: Oxford University Press), pp. 161-186.

- White, R.: 2010, 'You Just Believe that Because ...', *Philosophical Perspectives* 24(1), 573-615.
- Wielenberg, E. J.: 2010, 'On the Evolutionary Debunking of Morality', *Ethics* 120(3), 441-464.
- Williamson, J.: 1999, 'Countable Additivity and Subjective Probability', *British Journal for the Philosophy of Science* 50 (3), 401-416.
- Williamson, T.: 2006, 'Conceptual Truth' *Aristotelian Society Supplementary Volume* 80(1), 1-41.
- Zusne, L., and W. H. Jones.: 1989, *Anomalistic Psychology: A Study of Magical Thinking*, 2nd Edition.
Hillsdale, NJ: Lawrence Erlbaum.