

OUR RELIABILITY IS IN PRINCIPLE EXPLAINABLE

DAN BARAS

barasd@post.bgu.ac.il

ABSTRACT

Non-skeptical robust realists about normativity, mathematics, or any other domain of non-causal truths are committed to a correlation between their beliefs and non-causal, mind-independent facts. Harry Field and others have argued that if realists cannot explain this striking correlation, that is a strong reason to reject their theory. Some consider this argument, known as the Benacerraf–Field argument, as the strongest challenge to robust realism about mathematics (Field 1989, 2001), normativity (Enoch 2011), and even logic (Schechter 2010). In this article I offer two closely related accounts for the type of explanation needed in order to address Field’s challenge. I then argue that both accounts imply that the striking correlation to which robust realists are committed is explainable, thereby discharging Field’s challenge. Finally, I respond to some objections and end with a few unresolved worries.

I. THE BENACERRAF–FIELD CHALLENGE

At the core of mathematical Platonism is the commitment to there being mathematical truths that are not caused or constituted by our beliefs. Typically, Platonists also believe that mathematical truths are non-causal, meaning that they cannot cause anything (mathematical beliefs cause things, but they are supposed to be distinct from the mathematical truths), and that there are many mathematical truths that we, especially mathematicians, know. Harry Field (1989: 25–30), revising an earlier argument by Paul Benacerraf (1973), argues that the combination of these commitments gives rise to a detrimental argument against mathematical Platonism. Field’s argument generalizes to any theory with commitments similar to the above, namely any realism about a domain of non-causal truths. David Lewis (1986: 108) expresses concern about a parallel challenge to modal realism; David Enoch (2011: ch. 7) argues that Field’s challenge poses the most serious epistemological threat to robust normative realism;¹ and Joshua Schechter (2010) goes as far as to argue that Field’s argument threatens our knowledge of logic. In the remainder of this section, I will present the Benacerraf–Field challenge. In the sections that follow, I will explore a way of responding to the challenge.

¹ As some authors have noted (Enoch 2011: sec. 7.3; Setiya 2012; Berker 2014), Street’s evolutionary debunking argument incorporates several distinguishable lines of thought, one of which is an application of Field’s argument (Street 2006, 2008: 207).

The argument rests on an epistemic principle: If a theory implies that there is some apparently unexplainable striking correlation, that is a strong reason to reject the theory (Field 1989: 26). Considering some examples can help make this principle appealing. Suppose I flip what I take to be a fair coin 100 times and it lands in some insignificant sequence TTHTHHHHTTTT Despite the fact that the probability of it landing in this precise sequence is 2^{-100} , the sequence does not strike me as in need of explanation. The coin was bound to land in some sequence, and this sequence was as likely as any other. Compare this first coin example with another variant. Imagine I flip two presumably fair coins 100 times and, to my bewilderment, the coins land in perfectly matching sequences. Intuitively, the correlation between the coins is striking; we expect it must be somehow explainable. This is so despite the fact that the probability of the two coins landing in matching sequences is also 2^{-100} . Given that our background theory, which states that the coins and coin tosses are causally independent, fails to explain the striking correlation between the coin tosses, we have strong reason to reject that theory and seek an alternative. For instance, we may suspect that the coins are causally related after all. Similarly, using Field's (2001: 325) own example, suppose I believe that John and Judy have nothing going on between them. Suppose you now bring to my attention that throughout the past year, John and Judy were observed in close proximity to each other on countless occasions in varying locations: at the movie theater and various cafés, clubs, and museums. The observation seems like a good reason to seek an alternative theory because the massive correlation between the locations of the two of them over the past year seems too striking to leave unexplained. For instance, we may suspect that John and Judy are in a relationship.

The same principle seems to pose a challenge to robust realism about any domain of non-causal truths. If I am a non-skeptical normative realist, for instance, then I am committed to there being a massive correlation between my normative beliefs and the normative truths. Admittedly, in contrast to coins, we do not observe this correlation. We do not observe the truths, and then observe our beliefs, and find ourselves surprised by their striking correlation. In this case, it is my theory, rather than my observation, that commits me to there being this striking correlation, and that is enough to create a need for explanation. For robust realists the correlation seems hard to explain, because according to robust realism there is no causal or constitutive relation between the truths and our beliefs. Robust realists are committed to an apparently unexplainable striking correlation, and this seems like a strong reason to reject their theory. The argument can be summed up as a *reductio* argument as follows:

Field's Argument

Let D be a relevant domain of discourse (any discourse, such as normativity or mathematics, which appears to conform to robust realism).

- (1) *Non-Skeptical Robust Realism*: The robust realist has the following commitments:
- (a) *Realism*: There are D truths.
 - (b) *Cognitivism*: We have D beliefs.
 - (c) *Mind-Independence*: Our D beliefs do not cause or constitute the D truths.

- (d) *No-Causation*: The D truths do not cause or constitute our D beliefs.
- (e) *Non-Skepticism*: There is a correlation between the D truths and our D beliefs.

- (2) *Striking Correlation*: The correlation (e) is striking.
- (3) *Unexplainable Correlation*: If Non-Skeptical Robust Realism is correct then the correlation (e) is apparently unexplainable. ((c) & (d) rule out the most obvious kinds of explanation.)
- (4) *Epistemic Principle*: If a theory implies that there is an apparently unexplainable striking correlation, that is a strong reason to believe that the theory is false.

Therefore,

- (5) *Conclusion*: We have strong reason to reject Non-Skeptical Robust Realism (at least one of the commitments (a) through (e)).

There are varying ways in which realists can try to respond to the challenge. The Striking Correlation premise seems hard to deny. Some may question the Epistemic Principle; others may concede that Field's argument may count as a significant consideration against realism but argue that it is outweighed by other considerations.² Here I will explore one way of responding to the challenge, by rejecting the Unexplainable Correlation. My conclusion will be that if realists can explain, in the relevant sense of "explain," how they came to have the relevant set of beliefs, they can thereby explain how they came to be reliable regarding that set of beliefs, despite the fact that the relation between the beliefs and the truths will not be a causal relation. My conclusion is similar to that of some previous authors, specifically Joel Pust (2004) and more recently Justin Clarke-Doane (forthcoming a).³ In relation to these two authors, in this article I intend to make two significant contributions. First, my strategy in establishing my conclusion will be different from theirs. I will do so by exploring Field's Epistemic Principle more generally (section 3). In particular, I will suggest (a sketch of) an account of: when something calls for explanation, what sort of explanation does it call for? A second way in which my article goes further than those mentioned above is that (in section 4) I respond to an objection raised by Schechter (2010) that, to the best of my knowledge, has not been addressed previously.

2. WHAT KIND OF EXPLANATION IS CALLED FOR?

Clarke-Doane (forthcoming b) has recently complained that Field's formulation "is unclear at a crucial juncture. It is unclear what would count as an *explanation* in the relevant sense." While Clarke-Doane's focus is on the Benacerraf–Field challenge, the question he poses is quite general. Assuming that we accept Field's Epistemic Principle, we should want to know: When something calls for explanation, what sort of explanation

2 This is Setiya's (2012: ch. 2) strategy and it is also one of Clarke-Doane's strategies. Clarke-Doane (forthcoming a, sec. 4; forthcoming c, sec. 5) suggests that even if there is some relevant sense in which our reliability calls for explanation and seems unexplainable, so long as we are not given reason to believe that they are unsafe or insensitive, those beliefs remain justified. Clarke-Doane calls this principle Modal Security. Hence, another contribution of this article is that it attempts to respond to the Benacerraf–Field challenge without assuming Modal Security at the outset.

3 I thank an anonymous reviewer for referring me to these two articles, of which I was initially unaware. Regarding Clarke-Doane, see also the previous note.

does it call for? The easier it is to satisfy needs for explanation, the easier it might be for realists to respond to Field's challenge. On the other hand, the harder it is to satisfy such needs for explanation (i.e., if not just any explanation is sufficient, but some particular type of explanation is required), the harder it will be to respond to the challenge. For instance, Field himself, in later formulations of the argument, suggests that the type of explanation required is what he calls a unified explanation, and it is this particular type of explanation that he thinks mathematical Platonists must supply but cannot (Field 1996: 370). In this section, I will seek an answer to the general question by considering examples of the sort that motivated the Epistemic Principle to begin with.

As formulated above, the Epistemic Principle seems to imply that any sort of explanation will sufficiently satisfy a need for explanation. That cannot be correct. Consider our leading example of something that calls for explanation: You flip two coins 100 times and they land in perfectly matching sequences. The following proposition seems like a defective response to the challenge:

Coincidence: The two coins landed in matching sequences merely by coincidence.

Sometimes, saying that something happened by mere coincidence may be a good explanation. But, intuitively, when we come across striking correlations we have, *prima facie*, reason to believe that they are not coincidental. The fact that the above explanation is unsatisfactory demonstrates that when we have reason to believe that something is explainable, not just any kind of explanation will suffice. We have reason to expect a particular type of explanation. You may think that the Coincidence response does not count as an explanation at all and therefore does not demonstrate a need for a particular type of explanation. Consider, then, a different sort of explanation:

Initial Conditions and Mechanisms: On the first toss, I flipped the first coin at angle a_1 and velocity v_1 starting from the precise position s_1 , and these conditions and the laws of nature led the coin to land H. On the second toss, I flipped the first coin at angle a_2 and velocity v_2 starting from the precise position s_2 , and these conditions and the laws of nature led the coin to land T ... And so on for all 200 flips.

Arguably, there is a sense in which such a messy disjunctive specification of initial conditions and mechanisms counts as a genuine explanation. We can imagine a possible world in which something like the above explanation is the best possible explanation for the matching coin-toss sequences. If, as some have thought (Railton 1978, 1981), a causal explanation is just information about initial conditions and mechanisms that brought about the explanandum, such a long disjunctive explanation should suffice. Nevertheless, in the coin example, we have reason to believe that the above explanation is not the ultimate explanation for the correlation. Why? It still seems like too much of a coincidence that all these precise independent conditions obtained, which caused the coins to land in matching sequences.

In order to rule out one more prominent type of explanation, consider the following response:

Statistical Relevance: The coins are not perfectly symmetrical. Both coins are slightly weighted toward heads ($\approx 58.719\%$). Therefore, the probability of them landing in the same sequence is not $1/2^{100}$ but rather $20/2^{100}$.

Despite the statistical relevance of the proposed explanation – after all the probability of the explanandum is raised twenty times over – we would remain unsatisfied. Perhaps sometimes low statistical relevance is enough for something to count as an explanation (Salmon 1971), but when faced with something that calls out for explanation, we seem to have higher expectations.

The examples discussed above suggest that regarding some things, namely those that are striking, we have reason to believe that they are not the product of mere coincidence, or mere chance, improbable outcome, or whatever terms best capture what we mean. Rather, we seek an explanation that implies that the explanandum was highly probable. It seems that, regarding such explananda, we are satisfied with explanations that imply that the explanandum was highly probable, and dissatisfied with explanations that imply that the explanandum was improbable. Therefore, I suggest that the type of explanation needed to satisfy Field's Epistemic Principle is a high probability explanation, which can be characterized as follows:

High Probability Explanation: A high probability explanation consists of information about initial conditions that imply that the explanandum (i.e., the striking correlation) was highly probable.⁴

Given that the counterexamples above have eliminated the most salient competitors, and that the high probability account seems to fit well with our intuitive judgments of all the examples discussed thus far, it seems plausible that a high probability explanation is indeed the type required to satisfy a call for explanation.

Admittedly, the high probability account faces some challenges. First, presumably, as the example with the disjunctive explanation demonstrates, any explanandum will be highly probable given detailed enough information about initial conditions. Doesn't it therefore give the wrong prediction regarding highly disjunctive explanations as in the Initial Conditions and Mechanisms example above?

My response is that the problem with disjunctive explanations is that it remains too much of a coincidence that all the conjuncts came together in such a way as to produce the striking consequence. This judgment suggests that the problem with the disjunctive explanation is that the fact that those initial conditions obtained calls for explanation. In general, it is not enough that the explanation predicts the explanandum; the explanation itself has to be sufficiently probable. If the explanation is a long conjunction of apparently independent propositions, the probability of the conjunction can be expected to be extremely low. It seems plausible that this line of reasoning can generally explain why we expect a unified explanation. The fewer factors required to produce the striking result, the higher the probability that those factors will come together.⁵ Thus, Field's suggestion that we generally expect a unified explanation for striking phenomena does not

4 I do not have a worked-out account of the relevant interpretation of probability, but it seems the notion of probability relevant here is objective chance of the sort Lewis (1981) discusses and connects with subjective credences via his Principal Principle. It is an open question whether objective chance *just means* rational subjective credence prior to receiving what Lewis calls inadmissible evidence (like evidence that the coin in fact landed in a given sequence) or, alternatively, whether it is some property *in virtue of which*, prior to receiving inadmissible evidence, one would have reason to have a certain degree of belief.

5 White (2005b) suggests a similar explanation for why, *ceteris paribus*, simpler explanations are generally likelier.

contradict the high probability account but rather is predicted by that account. This point will resurface when I discuss Schechter's objection below.

Unfortunately, once we require that the explanation itself be highly probable as well, we set off a potentially infinite regression. You may worry that even if you go as far back as the big bang or cite the existence of God, it will still remain striking that there was a big bang or a God that brought about a coin landing 100 times heads. Hence, in order to be complete, an account of high probability explanations should be complemented with an account of where and why explanations come to an end. I have no such account to offer. It is obvious that explanations must come to an end somewhere, and hopefully, the analysis above is sufficient to motivate the thought that aside from those basic unexplained explanantia, high probability explanations are a plausible candidate for the type of explanation needed to satisfy Field's Epistemic Principle.

Roger White (2005a) suggests an alternative account, which I will examine alongside the high probability account. White suggests that when facing something that calls for explanation, we seek an explanation that implies that the explanandum will occur across a broad range of possibilities. For instance, if the coin is weighted toward landing heads, we can expect it to land heads over a wider range of possible velocities and angles. White calls this feature of explanations *stability*, which he defines in the following way:

Stability: An explanation of a fact F is stable to the extent that according to this explanation, F couldn't easily have failed to obtain. (White 2005a: 4)

White then explains his account in terms of nearby possible worlds: "F couldn't easily have failed to obtain" means that F obtains in all nearby possible worlds. White argues that we have reason to expect stable explanations for striking phenomena, and therefore stability is a good guide to the truth of an explanation.⁶ White's stability account is close to one of Field's suggestions, that the sort of explanation required is a counterfactually persistent explanation, which he initially introduces as "an explanation that worked in a general manner, independent of the details of where John happened to be each Sunday and where Judy happened to be" (Field 1996: 371). But White's formulation seems more plausible than Field's formulation. If the explanation implied that John and Judy could not have easily failed to be in the same place every Sunday for reasons that are dependent on the specific places in which they met, it seems we would be satisfied.

What is the precise relation between the high probability account and the stability account? That depends on what the precise relation between probability and modality is in general. I do not want to commit myself to any particular view on the matter. Fortunately, the argument in the next section can be formulated for both accounts.

I end this section by flagging a potential objection. One implication of both accounts is that the Epistemic Principle will never affect necessary facts (the necessary facts themselves, not the correlation between necessary facts and our contingent beliefs). If the idea behind Field's Epistemic Principle is that some things, namely the striking things, should be predicted by our theories, perhaps across a wide range of nearby possible worlds this worry will not arise for necessary facts, which are standardly assumed to

6 I note that in a more recent article, White (2007: 456) repeats this account of stability, but then it seems like the main argument of the article implies that stability is not a guide to explanation, contrary to the thesis of his earlier (White 2005a) article.

have probability 1 and to be true across all possible worlds. But one may worry that there can be necessary facts that are striking coincidences.⁷ If there were examples in which Field's principle intuitively applied to necessary facts (I have not found convincing examples), that would be a significant blow against the analysis of this section. I take this to be a genuine worry.

3. ALL THE REALIST NEEDS TO DO IS EXPLAIN HER BELIEFS

In this section I will argue that if either the high probability account or stability account is correct, then it is not so hard to respond to the Benacerraf–Field challenge after all.

I begin with a few preliminary remarks. It is important to note that the Epistemic Principle does not require that one actually explain the striking correlation, but rather that the striking correlation not be *apparently unexplainable*. Why? First, because normally, the thing to do when you come across a striking correlation is to seek an explanation, not to give up any of your beliefs. Only if you have reason to believe that no explanation exists will you have reason to revise some beliefs. In addition, the stronger principle, according to which one must actually have an explanation for any striking correlation, would have unwelcome consequences. Presumably, our ancestors were justified both in most of their perceptual beliefs and in their higher-order belief that their perceptual beliefs were correlated with the relevant truths, despite them having no idea how this correlation could be explained. The difference between perceptual beliefs and the problematic domains is that in the latter case, it is not only that realists lack an explanation but that, due to the lack of any causal relation, the correlation seems in principle unexplainable.

Notice in addition that the argument is a *reductio*. As such, it shows an undesirable consequence of *accepting* Non-Skeptical Robust Realism. Field goes even farther and is willing to grant that the realist has initial justification for her theory:

We start out by assuming the existence of mathematical entities that obey the standard mathematical theories; we grant also that there may be positive reasons for believing in those entities. (Field 1989: 25)

Although nothing in my discussion will rely on it, this is a good assumption to make. If there is no reason to believe that robust realism is true, then we need not bother with arguments against the view.

Given this setup, it is plausible to add one more assumption that most realists about non-causal truths accept, namely that non-causal truths, such as mathematical truths and basic normative truths, are necessary truths. This is a deep and controversial issue and I will not argue for it here.⁸ I am simply invoking a quite common way of thinking and adding the following to the imagined robust realist's package of commitments:

(f) *Necessity*: The basic D truths are necessary truths.

⁷ I thank Josh Schechter for presenting me with this challenge.

⁸ Field (1989: ch. 7) himself denies the necessity of mathematical truths. Rosen (unpublished ms.) has recently argued against the necessity of moral truths. Nothing in philosophy is uncontroversial. My argument should be understood as a conditional argument. If all these assumptions, which are broadly assumed by prominent writers in the debate, are correct, we will get a puzzling consequence.

Not all D truths need to be *basic* D truths. Take morality, for instance. There are many moral propositions that are obviously not necessary truths. It may be true in the actual world that it would be morally wrong for you to lie to your mother, but there are situations in which possibly it would be right to lie to your mother, such as when you don't want to worry her too much or when preparing a surprise party for her. Nevertheless, the common view is that there are basic moral truths (such as that lying is, *prima facie*, wrong) that are necessary truths. Moreover, it is commonly supposed that contingent moral truths are mere applications of the basic necessary truths. If this picture is correct, and if the realist can explain her reliability regarding the basic truths, it should not be too hard to explain her reliability about non-basic truths.

All the pieces are in place for the main argument of this section. Let us repeat the claims from above that will be used in the argument. First, for the sake of argument, we are assuming at the outset that the realist's beliefs are true. The question is only whether the correlation between the beliefs and the truths can be explained. Second, we are assuming that the relevant truths are necessary truths. Finally, we are assuming that the kind of explanation that the realist needs is a high probability explanation or a stable explanation.

I will argue that if all the above is correct, then in order to explain the correlation between her beliefs and the truths, all the realist needs to explain is her beliefs.

Let us first make the argument using the framework of probabilities. Suppose there is a correlation between A and B that calls for explanation (A and B are each a conjunction of a set of propositions). Normally – that is, when A and B are contingent – if you have a high probability explanation for A and a high probability explanation for B (E_1 and E_2 such that $P(A|E_1)$ and $P(B|E_2)$ are high), that does not imply that you have a high probability explanation of the conjunction $A \& B$. Why? Because if A and B are probabilistically independent, even if $P(A|E_1)$ and $P(B|E_2)$ are high, that does not imply that $P(A \& B|E_1 \& E_2)$ is high. To demonstrate using one of our leading examples: We may have individual explanations for why John and Judy went to every place in which they were spotted throughout the year:

On March 1 John went to a café because he needed a break from working on an important draft. On March 9 John went to a movie in order to celebrate finishing the draft. On March 1 Judy went to a café because she was in the mood for a muffin. On March 9 Judy went to a movie because it was her last chance to see a movie that she really wanted to see ...

But we do not thereby have an explanation for why throughout the year John and Judy were repeatedly spotted together. Even given all of these explanations, the probability of the conjunction of places and times in which John and Judy were spotted together, and even the probability of the correlation between John's and Judy's hangouts, remains exceedingly low.

Things are different when B is a conjunction of necessary facts (i.e., probability 1).⁹ If $P(B) = 1$, on the standard Kolmogorovian framework, it follows that $P(A \& B|E) = P(A|E)$. In other words, any explanation that entails high probability of A will also entail high

9 Arguably, under an epistemic/subjective interpretation of probability, it can be rational to have less than credence 1 even for logically necessary propositions. This is why the logical omniscience implied by the standard Bayesian framework is a problematic idealization. Following my note earlier (fn. 4), the probability function here is supposed to represent not one's current ideal state of beliefs but rather what one should have expected the world to be like given our background theories.

probability of the conjunction of A and B. This combines with another feature of probability 1 propositions. We may ask, what is the probability that I will have a true belief about some proposition p given explanation E, which contains information about initial conditions and laws of nature? If p is a necessary truth, meaning that p can only be true, the probability that I have a true belief about p equals the probability that I believe p ($P((p \text{ and I believe } p) \text{ or } (\sim p \text{ and I believe } \sim p)|E) = P(\text{I believe } p|E)$). From the above it follows that for any necessary truth or conjunction of necessary truths, if I have a high-probability explanation for my beliefs, I thereby have a high-probability explanation for the correlation between my beliefs and the truths.

The same idea can be demonstrated using the possible-worlds framework and White's notion of a stable explanation. Normally, if we want to provide a stable explanation of a correlation between contingent sets of facts A and B, we want to explain why in most nearby worlds, wherever A remains the same, so does B, and wherever A changes, B changes accordingly. Things are different when one of the sets of facts is stable across all nearby worlds. Suppose we flip two supposedly fair coins and they land in a perfectly matching sequence. We agreed above that this would be something that calls for explanation. Suppose now we learn that the first coin was double headed, and the reason that it landed 100 times heads is that it could not have landed otherwise. Once we have this information, all that remains to be explained is why the second coin landed 100 times heads as well. The explanation of the second coin need not have any resemblance to the explanation of the first; we can be satisfied, for instance, with the hypothesis that the second coin is not double headed but rather is weighted toward landing heads.

For contingent truths, a stable explanation for my belief that p, which is actually true, is not necessarily a stable explanation for my having a true belief about p. Even if I could not have easily disbelieved p, if p easily could have been false, I easily could have had a false belief about p. For example, suppose I believe that *The Philosopher's Annual* will choose my article as one of the top ten articles of the year. Suppose further that my article will indeed appear in *The Philosopher's Annual*. Finally, suppose that the best explanation for this belief of mine is that it is the product of wishful thinking. The explanation of my belief seems like a stable explanation. Due to wishful thinking I would believe that my work would appear in the annual across a wide range of possible worlds. But I do not have a stable explanation for the correlation between my belief and the truth, because the editors easily could have not noticed my article and, due to wishful thinking, I would falsely believe that my work will be chosen as one of the top ten.

Things are different with regard to explanations of beliefs in necessary facts. If p is a necessary fact, any stable explanation of my belief that p will also be a stable explanation for my belief in the truth about p, because the truth about p is just p in all possible worlds. This point about a single belief applies just as well to sets of beliefs. Hence, for any set of necessary truths that I believe, a stable explanation for my beliefs will be a stable explanation for the correlation between my beliefs and the truths in that domain. In other words, any explanation that implies that I could not easily have had basic beliefs that are different from the ones I actually have, assuming those beliefs are true, will also explain why I could not easily have had false beliefs.

So here's the main claim of this section:

In order to explain ("explain" in the sense relevant to the Epistemic Principle) the correlation between her beliefs and any set of necessary truths, all the realist will need to do is explain her beliefs.

If this is correct, the striking correlation no longer seems unexplainable. Even if we do not yet have the ultimate explanation of why we have the basic moral or mathematical beliefs that we have, we have good reason to believe that one is forthcoming, whether it be evolutionary or otherwise. We need not dirty our philosophical hands with the details of the explanation; we can leave that to experts in the empirical sciences. The realist can rest assured that as long as the explanation of her beliefs is a high probability (or stable) explanation, it will suffice as a response to Field's challenge.

4. SCHECHTER'S OBJECTION

In this section, I will respond to an interesting challenge raised by Josh Schechter. Discussing Schechter's objection will give me the opportunity to make some additional clarifications. For brevity, I will focus on the high probability account, but similar considerations apply to the stability account as well.

Schechter has us consider the following example:

Suppose it were true that every day in March, the number of people who took the New York subway was a prime number. Suppose we could provide an elaborate explanation of why various people did (or did not) take the subway on particular days in March. Even though this explanation would entail that a prime number of people took the subway each day in March, it would not explain this fact. The fact that there was always a prime number of people would remain mysterious. (Schechter 2010: 447)

The example is supposed to demonstrate that explanations are not generally closed under logical consequence. In the example, the thought is that even if you can explain why each and every person did or did not get on the subway each day, that will not suffice as an explanation for the fact that the numbers were prime, despite the latter being a necessary consequence of the former. Likewise, you may suspect that even if we can explain why we have the beliefs that we do, and those beliefs are necessarily true, it does not follow that we can thereby explain the correlation between our beliefs and the truths. Thus goes Schechter's line of thought.

The first problem with Schechter's analysis is that there is an alternative explanation for our intuition that the given disjunctive explanation is insufficient, one that is perfectly consistent with the argument of this article. Think again about the example. You are probably thinking: I understand why Aaron got on the train on March 1; his car broke down that day. I understand why Dalia did not get on the train on March 1; she was in the midst of a three-day conference out of town. This continues for each and every potential subway passenger. But what were the chances that Aaron's car would break down that day, and that Dalia would be in a conference that day, and all of the massive list of explanantia, such that the sum of people who got on the subway on March 1 would end up being prime, not to mention the rest of the days of the month? The chances remain very low.

More generally, we should distinguish between two different principles¹⁰:

Agglomeration of Explanations: If I have a high probability explanation for p and a high probability explanation for q , I thereby have a high probability explanation for $p \& q$.

Closure of Explanations: If I have a high probability explanation for p , I thereby have a high probability explanation for any necessary consequence of p .

¹⁰ Thanks to Ofer Malcai for helping me see this.

The probabilistic analysis of explanations does not imply the agglomeration principle because, unless the probability of one of the conjuncts is 1, the probability of the conjunction will be lower than the probability of each individual conjunct. The longer the conjunction, the lower the probability of the conjunction, even when the probability of each individual conjunct, given our background theories, is high. Now, the fact that the number of passengers each day of the month was prime is a necessary consequence of the conjunction “Aaron, Betty, Charlie ... were on the subway on March 1. Dalia, Ethan, Fiona ... were not on the subway on March 1 ... (and so on for each day of the month)”; it is not a necessary consequence of any one of the individual conjuncts. In the subway example we are likely imagining a probabilistic explanation of why each person did or did not get on the subway every day such that the conjunction of the explananda has an extremely low probability given the disjunctive explanans. Given such a disjunctive explanation, the probability of what is implied by the conjunction, that the sums throughout March were all prime numbers, is extremely low as well. Hence, we have a clear intuition that even if I can explain why each person did or did not get on the subway each day, I cannot thereby explain why the sums were prime numbers. But this intuition is well explained by the probabilistic considerations that imply the falsity of *agglomeration* and therefore provides no evidence for the claim that *closure* is false.

In order to test Schechter’s claim, we need to imagine a case in which we have an explanation that implies that the conjunction “Aaron, Betty, Charlie ... were on the subway on March 1. Dalia, Ethan, Fiona ... were not on the subway on March 1 ... (and so on for each day of the month),” not just each individual conjunct, was highly probable. Such a situation is so hard to imagine that I doubt we have any reliable intuitions about it. Indeed, if we were in such a position, I concede that the fact that the sums are prime numbers would remain mysterious in some sense. But I doubt it would be in any sense epistemically significant. To take a similar example, the existence of a world with conscious beings like us will probably remain mysterious in some (psychological?) sense even after future scientists and philosophers have the best explanation possible. Given that they will have the best explanation possible, they will obviously have no reason to doubt any of their beliefs. The lesson is that not every sort of mystery ought to give rise to a Benacerraf–Field challenge. Hence, once we adjust Schechter’s example such that the probability of the conjunction is high, even if the prime numbers remain mysterious, it is far from clear that there will be a need for explanation in Field’s sense, one that requires a revision of beliefs.

A final remark on Schechter: Schechter’s argument was supposed to be for the claim that explanations are not closed under logical consequence. Now you may worry that there are plenty of examples to demonstrate that indeed explanations are not closed under logical consequence. For instance, an explanation of why Adam is a bachelor does not necessarily entail an explanation of why Adam is a male. But such examples do not carry over to the debate at hand. Even if explanations in general are not closed under logical consequence, in the context of Field’s challenge, I have argued that the need is for a specific type of explanation, a high probability or stable explanation given our background theories. The question is whether the sort of explanation the Epistemic Principle requires is or is not closed under logical consequence. In ordinary conversation, an explanation of why Adam is a bachelor does not necessarily entail an explanation of why Adam is a male. That is because in ordinary conversation pragmatic considerations determine what sort of explanation would be appropriate to give. When one requests an explanation for why Adam is a bachelor, we are normally allowed to take for granted

many facts, such as Adam’s masculinity, and only have to explain why he is a bachelor as opposed to being married. The interests of those requesting explanations can intuitively explain the pragmatics of explanations. On the other hand, in the context of debates on ontological theories, it would be puzzling if such pragmatics should play an epistemic role in determining which theories should be rejected because they lack an explanation to a striking correlation. In our case, it would seem weird if there were some fact we were permitted to take for granted when explaining why each passenger did or did not get on the train each day, but not when explaining why the sum was prime.

To further strengthen this point, I draw an important distinction from Michael Strevens (who attributes it to Bromberger (1965)) between two senses of explanation:

Philosophers sometimes talk as if an explanation were something out in the world, a set of facts to be discovered, and sometimes as if it were a communicative act. (Strevens 2008: 5)

As Strevens goes on to explain, explanation as a communicative act is subject to the same sorts of pragmatic constraints that any communicative act is subject to. Strevens’s distinction is helpful in the context of Field’s challenge. Pragmatics that disrupt the closure of explanations are related to explanation as a linguistic phenomenon. In the context of Field’s challenge, in which we are concerned with causal explanations provided by our background theories, it is the ontological meaning of explanation we are using, and it is plausible to assume that such explanations are closed under logical implication, as the high probability and stability accounts imply. Hence, examples like Adam the bachelor are irrelevant to the debate at hand.

5. PRE-ESTABLISHED HARMONY EXPLANATIONS

In this section, I will point out an interesting implication of this article’s thesis with regard to realism about normativity. In defense of normative realism against the Benacerraf–Field challenge, several writers have suggested what I will call, following David Enoch, pre-established harmony explanations (PEH).¹¹ The general thought is that even without a causal or constitutive relation between our beliefs and the truths, we can explain why a PEH between them is to be expected.

The core of Enoch’s explanation is expressed in the following passage:

Selective forces have shaped our normative judgments and beliefs, with the “aim” of survival or reproductive success in mind (so to speak). But given that these are by and large good aims – aims that normative truths recommend – our normative beliefs have developed to be at least somewhat in line with the normative truths. (Enoch 2011: 168)¹²

11 It is noteworthy that prominent responses in the domain of mathematics are hard to apply to realism about normativity. Among them are: (i) All of the infinite possible consistent mathematical systems are factual and therefore our reliability in logic is sufficient to explain our reliability in mathematics (Balaguer 1995). (ii) All mathematical truths can be reduced to a single, basic, set theoretic truth and therefore it is sufficient to explain our knowledge of this single truth to explain our mathematical knowledge (Burgess and Rosen 1999: 45). (iii) Mathematical reliability is adaptive and therefore evolutionary theory can explain our reliability about mathematics (Schechter 2013). Whether or not any of these responses is plausible is not a question I can address here.

12 Nozick (1981: 346) previously suggested this sort of explanation in response to a different version of an explanatory challenge; Parfit adopts Nozick’s solution (Parfit 2011: sec. 114), and Berker (2014)

According to Enoch, even though there is no causal link between our moral beliefs and the moral facts, a conjunction of metaphysical facts and natural facts explains why the striking correlation is not coincidental. Given that we are the products of evolution by natural selection, it is to be expected that our cognitive capacities would be constructed such that we would be inclined to believe that actions that enhance our survival are good. Given that our survival is, *prima facie*, good, it is to be expected, generally speaking, that things that promote our survival are good.¹³ So our theories have sufficiently predicted that we have evolved to have many true normative beliefs.

The hypothesis I have explored in this article is that the PEH strategy is easy to generalize – so easy that any explanation of our beliefs can generate a PEH explanation of the correlation between our beliefs and the truths. As it turns out, this hypothesis is not quite accurate. Not just any explanation of our beliefs will do. The main claim of this article has been that any high probability (or stable) explanation of our beliefs, evolutionary or otherwise, will do. This may come as a surprise to some people. I have discovered that many colleagues worry that Enoch’s PEH explanation depends very much on the details and that it is far from clear that once the details of our best moral theory and evolutionary theory are specified, Enoch’s strategy will turn out plausible. But if I am right, Enoch need not worry about the details. Any evolutionary explanation will do, no matter what your basic first-order normative commitments may be.¹⁴

6. SOMETHING HERE SMELLS FISHY

You may think that something here smells fishy. Field’s argument initially seemed powerful. Admittedly, sometimes challenges seem appealing but upon closer examination, the appeal disappears. Some may feel that this is what happened here. Others, myself in some moods included, might feel that if indeed the challenge is as weak as this article implies that it is, something must have gone wrong. Either we have misunderstood Field’s argument or something is wrong with the response. I will end this article by pointing out a few of my own suspicions.

Perhaps the strongest suspicion is that in assuming that her beliefs are true, the realist is illegitimately begging the question. In general, it is appealing to think that when considering

analyzes it in a helpful manner. Evolutionary explanations have also been used to respond to Field’s challenge as applied to mathematics (Clarke-Doane [forthcoming b](#)) and logic (Schechter 2013). Alternative PEH explanations of our reliability about morality are suggested by Skarsaune (2009) and Wielenberg (2010).

- 13 That is not to say that survival is always good at any price, nor is it to say that there cannot be forms of goodness that are independent of survival.
- 14 Several colleagues, including Enoch himself, have expressed in private correspondences their intuition that an essential feature of Enoch’s explanation is that we are the products of a process with an intuitively good “aim.” Now, I’m not sure what it means for nature to have an “aim.” Perhaps all it means for nature to have an aim *x* is that nature is such that it will likely bring about *x*. One problem with this suggestion is that even if this intuition is correct, so long as it turns out that nature had the aim of producing creatures with true moral beliefs, it seems I should believe that nature had a good aim. Now, no matter what my moral beliefs are, if they are justified, I should believe that if nature “aimed” at providing me with *those* beliefs, then nature aimed at providing me with (what I take to be) *the true beliefs*. That seems like a very good aim. Therefore, even if we required such an additional condition, it seems like it will be satisfied whenever I have a high probability explanation of my beliefs, and the claim of this article remains intact.

whether p , one ought to set aside one's belief that p , and perhaps even some of one's reasoning for p , and only take into consideration *independent* support for p . Similarly, in the context of questioning our beliefs, it would be irrational to assume their truth, not to mention their modality. Principles of independence have been at the heart of the large debate on higher-order evidence (Christensen 2011). As appealing as it may seem initially, simplistic formulations imply an intolerable global skepticism (White 2010). Perhaps, then, the argument of this article is a good response to Field's argument, but there is some stronger argument in the vicinity that will include a reasonable principle of independence.

A different worry is that neither the probabilistic nor the modal account of explanation is the correct way of understanding what sort of explanation is needed to answer Field's challenge. Despite my efforts to support these accounts using examples and diffuse some potential counterarguments, there are still some reasons to remain suspicious. First, I rely on a very vague notion of objective probability. I know of no account of objective probability without problems, and therefore I worry that the vague notion of probability that I rely on will turn out unreasonable or incoherent. Second, as discussed earlier, it is hard to come up with applications of Field's Epistemic Principle to necessary facts. Yet the mere fact that it seems plausible that necessary facts can be striking coincidences in some sense, probably epistemic, gives reason to suspect that maybe the notion of probability assumed in this article is inappropriate. Third, the fact that I have no account of when explanations come to an end in this context worries me. For any set of initial conditions and laws of nature, you can always wonder whether it is not too big a coincidence that those were the precise initial conditions and laws of nature such that the thing to be explained ends up being highly probable.

Finally, for some domains, I worry that we should not expect a high probability explanation of our beliefs to be forthcoming. It is well known that there are vast and deep moral disagreements among people and even within the same person over time. Obviously there is a sense in which it is a coincidence that I have the precise moral beliefs that I do rather than, say, the moral beliefs that you do, or that my ancestors did. This gives us reason to believe that whatever explains my moral beliefs will not necessarily predict my precise set of beliefs. If no high probability explanation is to be expected, neither should a high probability explanation of the correlation between my beliefs and the truths be expected. So there is a worry that the realist will not be able to respond to Field's challenge by rejecting the No Explanation premise after all.^{15,16}

REFERENCES

- Balaguer, M. 1995. 'A Platonist Epistemology.' *Synthese*, 103: 303–25.
 Benacerraf, P. 1973. 'Mathematical Truth.' *Journal of Philosophy*, 70: 661–79.

15 I thank Ron Aboodi for raising this objection.

16 Although my name alone stands beside the title of this paper, it is better viewed as the product of a collaboration with many colleagues who have immensely contributed to my work by providing me with challenges and constructive comments throughout the long writing process. I am especially in debt to Ron Aboodi, Michael Antony, Selim Berker, David Enoch, Eran Fish, Yehuda Gellman, Ran Lanzet, Arnon Levy, Ofer Malcai, Eli Pitcovski, Ronni Sadovski, Said Saillant, Josh Schechter, Katia Vavova, Roger White, an anonymous referee and participants at my presentations at the Underground Workshop at the Hebrew University, and MATTI at MIT.

- Berker, S. 2014. 'Does Evolutionary Psychology Show That Normativity Is Mind-Dependent?' In J. D'Arms and D. Jacobson (eds), *Moral Psychology and Human Agency: Philosophical Essays on the Science of Ethics*, pp. 215–52. Oxford: Oxford University Press.
- Bromberger, S. 1965. 'An Approach to Explanation.' *Analytical Philosophy*, 2: 72–105.
- Burgess, J. P. and Rosen, G. 1999. *A Subject With No Object*. Oxford: Oxford University Press.
- Christensen, D. 2011. 'Disagreement, Question-Begging and Epistemic Self-Criticism.' *Philosophers' Imprint* 11 (6).
- Clarke-Doane, J. Forthcoming a. 'Debunking and Dispensability.' In N. Sinclair and U. Leibowitz (eds), *Explanation in Ethics and Mathematics*. Oxford: Oxford University Press.
- Forthcoming b. 'Debunking Arguments: Mathematics, Logic, and Modal Security.' In R. Richards and M. Ruse (eds), *Cambridge Companion to Evolutionary Ethics*. Cambridge: Cambridge University Press.
- Forthcoming c. 'What Is the Benacerraf Problem?' In F. Pataut (ed.), *New Perspectives on the Philosophy of Paul Benacerraf: Truth, Objects, Infinity*. Dordrecht: Springer.
- Enoch, D. 2011. *Taking Morality Seriously*. Oxford: Oxford University Press.
- Field, H. 1989. *Realism, Mathematics, and Modality*. Oxford: Blackwell.
- 1996. 'The A Priority of Logic.' *Proceedings of the Aristotelian Society*, 96: 359–79.
- 2001. *Truth and the Absence of Fact*. Oxford: Oxford University Press.
- Lewis, D. 1981. 'A Subjectivist's Guide to Objective Chance.' In W. L. Harper, R. Stalnaker, and G. Pearce (eds), *Ifs*, pp. 267–97. The University of Western Ontario Series in Philosophy of Science. Amsterdam: Springer Netherlands.
- 1986. *On the Plurality of Worlds*. Oxford: Blackwell.
- Nozick, R. 1981. *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Parfit, D. 2011. *On What Matters*. Oxford: Oxford University Press.
- Pust, J. 2004. 'On Explaining Knowledge of Necessity.' *Dialectica*, 58: 71–87.
- Railton, P. 1978. 'A Deductive-Nomological Model of Probabilistic Explanation.' *Philosophy of Science*, 45: 206–26.
- 1981. 'Probability, Explanation, and Information.' *Synthese*, 48: 233–56.
- Rosen, G. 'What Is Normative Necessity?' Retrieved from https://www.academia.edu/9159728/Normative_Necessity. Manuscript.
- Salmon, W. C. 1971. *Statistical Explanation and Statistical Relevance*. Pittsburgh, PA: University of Pittsburgh Press.
- Schechter, J. 2010. 'The Reliability Challenge and the Epistemology of Logic.' *Philosophical Perspectives*, 24: 437–64.
- 2013. 'Could Evolution Explain Our Reliability about Logic?' *Oxford Studies in Epistemology*, 4: 214–39.
- Setiya, K. 2012. *Knowing Right From Wrong*. Oxford: Oxford University Press.
- Skarsaune, K. O. 2009. 'Darwin and Moral Realism: Survival of the Fittest.' *Philosophical Studies*, 152: 229–43.
- Street, S. 2006. 'A Darwinian Dilemma for Realist Theories of Value.' *Philosophical Studies*, 127: 109–66.
- 2008. 'Reply to Copp: Naturalism, Normativity, and the Varieties of Realism Worth Worrying about.' *Philosophical Issues*, 18: 207–28.
- Strevens, M. 2008. *Depth: An Account of Scientific Explanation*. Cambridge, MA: Harvard University Press.
- White, R. 2005a. 'Explanation as a Guide to Induction.' *Philosophers' Imprint*, 5: 1–29.
- 2005b. 'Why Favour Simplicity?' *Analysis*, 65: 205–10.
- 2007. 'Does Origins of Life Research Rest on a Mistake?' *Noûs*, 41: 453–77.
- 2010. 'You Just Believe That Because ...' *Philosophical Perspectives*, 24: 573–615.
- Wielenberg, E. J. 2010. 'On the Evolutionary Debunking of Morality.' *Ethics*, 120: 441–64.

DAN BARAS is a PhD candidate at Ben Gurion University of the Negev. His primary fields of research are epistemology and metaethics.
