

Extreme Betting

Javier González de Prado Salas

IFILNOVA, Universidade Nova de Lisboa

jgonzalezdeprado@gmail.com

**This is a penultimate draft of an article to appear in Ratio (doi:10.1111/rati.12217).*

Please refer to the published version.

Abstract: It is often thought that bets on the truth of known propositions become irrational if the losing costs are high enough. This is typically taken to count against the view that knowledge involves assigning credence 1. I argue that the irrationality of such extreme bets can be explained by considering the interactions between the agent and the bookmaker. More specifically, the agent's epistemic perspective is altered by the fact that the bookmaker proposes that unusual type of bet. Among other things, being willing to offer a bet with unnecessarily harmful losing costs is likely to undermine the baseline level of trustworthiness required for it to be rational to engage in betting exchanges. This sort of explanation does not require granting either that we assign credence lower than 1 to known propositions or that knowledge is sensitive to practical stakes. Moreover, I show that, in our ordinary lives, we frequently perform actions that we know would be disastrous if certain conditions did not obtain. This behaviour can be seen as a form of implicit extreme betting and, nevertheless, it is often rational.

1. INTRODUCTION

What I will call an *extreme bet* – following Brown (2008) – is a bet on the truth of p such that the costs of losing are very high (say, you will be enslaved forever) and the winning rewards are comparatively low. Extreme bets are typically regarded as a test case for theories that hold that knowledge involves assigning credence 1 to the known

proposition.¹ According to standard decision theory, it should be irrational to reject an extreme bet on the truth of a proposition to which one assigns credence 1, as long as winning the bet has some positive utility and regardless of the losing costs. However, there is the intuition that extreme bets on the truth of known propositions are not always rationally acceptable: if the costs of losing are very high, it will eventually become rational to reject the bet, even when dealing with trivial, seemingly secure pieces of knowledge (Brown, 2008, pp. 175-176; Clarke, 2013; Greco, 2013; Hawthorne, 2004; Hawthorne & Stanley, 2008, pp. 587-589; Hawthorne, 2004; Maher 1993, p. 133; Wedgwood, 2012; Williamson, 2000; Worsnip, 2016, p. 554).

In this vein, Hawthorne (2004, p. 29 fn. 72) claims that ‘I wouldn’t even bet on the law of noncontradiction at any odds, and I think myself rational on that score.’ Williamson (2000, p. 213) makes a similar point: ‘If one’s credence in P is 1, one should be willing to accept a bet on which one gains a penny if P is true and is tortured horribly to death if P is false. Few propositions pass that test.’ Hawthorne and Williamson give voice in these passages to the generalized view that extreme betting, when extreme enough, is irrational in relation to any (or almost any) proposition (or, at least, that it is rational to refrain from accepting such bets).²

On the face of it, this intuition about extreme bets has important epistemological implications. One possible reaction is to consider that the irrationality of extreme betting motivates scepticism: extreme bets would be taken to show that we actually lack knowledge about most ordinary, trivial propositions. An alternative response is to argue that knowledge is sometimes compatible with some forms of uncertainty and doubting behaviour (Brown, 2008; Reed, 2010). A further possibility is to maintain that degrees of credence are sensitive to practical factors such as how high stakes are (Clarke, 2013; Greco, 2015; also Stanley, 2005).

¹ The view that knowledge involves maximal credence is motivated, among other considerations, by lottery cases and, more generally, by the intuitive plausibility of multi-premise closure for knowledge (see for instance Hawthorne, 2004). Some authors also defend that outright belief requires maximal credence (Clarke, 2013; Dodd, 2017; Greco, 2015).

² Of course, there may be all sorts of reasons that make it rational to reject betting, even when the bet is guaranteed to yield a reward. For instance, one may be in a rush, or one may have some psychological aversion to the very idea of betting. Assume that in the examples considered here the utility of winning the bet is high enough to compensate for these possible costs.

My aim here is not to take a stand on the merits of these responses to our intuitions about extreme bets, but rather to argue that such intuitions do not have the epistemological significance that is usually attributed to them. In Section 2 I show that the examples in which it is rational to reject extreme bets can be accounted for by appeal to considerations that are largely neutral about the epistemological debates in question. In particular, we should take into account how the perspective of the agent is modified by the fact that she is being proposed an extreme bet by a bookmaker. I will argue that agents generally have good reasons not to trust the reactions and motives of a bookmaker offering an extreme bet, given the terrible losing costs that such a bookmaker is willing to enforce – and it may be rational for an agent not to engage in some betting interaction if she is not in a position to predict and rely on the reactions of the bookmaker. Furthermore, in Section 3 I show that it is possible to think of many cases of extreme betting (implicit in our everyday agency) in which these considerations do not apply and therefore it is rational to take the bet (and irrational to reject it). As it turns out, extreme bets are not always irrational, and when they are it is because of considerations that have no direct bearing on the thesis that knowledge involves maximal credence.

It should be stressed that intuitions about extreme betting play an important role in a number of debates, for instance debates about the norms of practical reasoning (Brown, 2008; Hawthorne & Stanley, 2008), about the relations between belief and credences (Clarke, 2013; Dodd, 2017; Greco, 2015), about pragmatic encroachment on knowledge (Fantl & McGrath, 2009; Stanley, 2005) or about decision theory (Greco, 2013; Williamson, 2005). It is not my goal to defend specific positions on these issues. However, the conclusions of this paper will have implications for the dialectics of such debates, insofar as the argumentative force of appeals to extreme betting will have to be reassessed.

2. IRRATIONAL EXTREME BETTING

Think of David, who, like most of us, has no doubts about which city his house is in: he is certain that his house is in London. One day David meets Mary who proposes to him the following bet: if David is right about which city his house is in, he will be rewarded

with a piece of candy;³ but if he is mistaken, he and all his friends will be enslaved and tortured for as long as they live. Mary seems to be serious and able to enforce the bet.

Arguably, it would be unwise of David to accept this extreme bet. Does this mean that he does not actually know whether his house is in London? If so, is it because he never actually knew for certain that his house was in London, or rather because his knowledge is sensitive to how high stakes are? I will argue that the rationality of rejecting the extreme bet can be explained without answering these questions, but focusing instead on how David's epistemic perspective is distorted by the fact that such an unusual bet has been proposed to him.

2.1. The untrustworthy character of extreme bookmakers

One first thing to say is that, if Mary is not joking and is able and willing to enforce the bet, she is certainly a dangerous and inordinate character (she seems ready to enslave and torture random people on a whim). Moreover, her motives are very difficult to understand: what are her reasons for proposing this weird, perverse bet? The unintelligibility of her behaviour makes her unpredictable and difficult to trust. Who knows if she will actually keep her word? Perhaps she will try to harm David even if he wins the bet. One should probably avoid someone that would propose and make effective this atrocious bet. There are surely good reasons to limit one's interactions with this type of person and to stay away from her. In particular, there are very good practical reasons not to agree to take a bet proposed by a menacing individual like this – reasons that have nothing to do with David's certainty about where his house is. The rationality of agreeing with Mary to take the bet does not only depend on the odds of the bet, but also on what is reasonable for David to expect about Mary's reactions.⁴ Can she be trusted to honour their agreement? Will she behave violently if she loses? David may have justified doubts about the answers to these questions, given that Mary's extraordinary betting proposal makes her appear as a dubious, hard to read individual.

Thus, the radical potential consequences of the extreme bet undermine the minimal presuppositions of trustworthiness underpinning the rationality of betting behaviour. In general, it is rational to avoid high stakes betting exchanges if one is not in a position to rely on the bookmaker and to anticipate the ways in which she may react

³ Or any other reward whose utility is just high enough to outweigh the possible costs of taking the time to consider the bet.

⁴ For discussion, see Greco (2013).

(e.g. if one does not know whether the bookmaker will keep to the terms of the bet). Arguably, this is what happens in cases like David's. Regardless of his degree of certainty about the relevant proposition, David has good reasons to be suspicious about the character and intentions of Mary the bookmaker, and this may make it rational for him to reject taking the bet.

2.2. Extreme bets as an additional source of evidence

In addition to there being reasons for David to be wary of Mary's possible reactions and motivations, facing such an unusual bet changes David's evidential situation, as Ian Hacking already noted (Hacking, 1965, pp. 206-207; also Dodd, 2017; Eriksson & Rabinowicz, 2013).⁵ The fact that Mary is willing to propose this weird bet must be counted among the evidence available to David when deciding what to do. In particular, this fact may give David reasons to think that Mary has information that he lacks (see Hawthorne, 2004, p. 176, fn. 36). Perhaps she is aware of some extraordinary turn of events that made it false that David's house is in London (say, there was some so far unnoticed problem with the house's documents, with the result that David does not actually own the London house; or maybe the house was somehow transported to other city). Why may someone propose a bet with such ruthless losing costs (and no apparent gain)? One of the few rational answers is that Mary wants to harm David and is laying a trap for him. She may be manipulating the situation in such a way that David turns out to be wrong and therefore would lose the bet. After all, someone capable of enforcing the radical consequences of losing the bet (i.e. enslaving and torturing David and his friends forever) may be capable as well of other extraordinary feats – such as transporting David's house to another city without him noticing. This possibility becomes still more relevant in fanciful versions of the example, in which Mary is a magic genie (that is why she has the power to enforce the bet). Once a magic genie – or a quasi-omnipotent psychopathic bookmaker – comes into the picture, one would start doubting many things that are typically taken for granted. Such doubts will be reflected by a reluctance to accept the bet.

⁵ Eriksson and Rabinowicz (2013) discuss several scenarios in which the agent's degree of credence in p differs from her degree of credence in p conditional on the supposition that she is proposed a bet on p . The agent's disposition to bet on p would depend on this latter conditional credence.

So, the unexpectedness of the bet may make it reasonable for David to contemplate possibilities that would be usually discarded as outlandish. Notice that what happens here is that David acquires new evidence – evidence giving him reasons to have doubts that he would not otherwise have. So, a theory of knowledge as involving credence 1 would be able to account for the irrationality of accepting Mary's bet just by claiming that knowledge and degrees of credence depend on the evidence available to the subject. The idea would be that the variation in David's credence between ordinary circumstances and a special situation where Mary proposes the bet is explained by the fact that he has different evidence in each case.⁶ Nothing in this requires granting that knowledge or credences are sensitive to practical stakes, or conceding that David does not have credence 1 when the bet is not proposed to him (note that the fact that reasonable doubts arise when David faces a bookmaker proposing an extreme bet does not mean that he must have the same doubts in ordinary circumstances).⁷

It may be argued that there are cases in which the agent's evidence is not relevantly altered by the proposal of the extreme bet. Perhaps Mary the bookmaker lets David choose the proposition the extreme bet will be about, so that there are no suspicions that Mary is selecting a special proposition about which she has privileged information (this scenario is suggested by Hacking, 1965, p. 207). In this case, it seems that David would have no reasons to think that Mary is aware of extraordinary circumstances affecting the truth of the chosen proposition – after all, Mary could not know in advance which proposition David was going to choose.⁸ I have to say that, in this sort of case, I find the intuition that it is irrational to accept the extreme bet less strong. At any rate, the considerations at the beginning of the section would still apply. Mary the bookmaker would be someone with the will and means to execute the terrible punishment resulting from David's losing the bet, with no apparent gain on her side. As

⁶ Of course, if we want to claim that a degree of credence 1 in a proposition may be lowered after new evidence becomes available, it will be necessary to modify standard Bayesian models of belief updating (see Clarke, 2013, pp. 11-13; also Williamson, 2000, ch. 10).

⁷ To be clear, the change in credences would not be due to the change in practical stakes, but to the change in David's evidence. I remain neutral about whether credences can be affected directly by shifts in practical stakes. My point is that it is not necessary to accept such pragmatic encroachment on credences in order to deal with examples in which extreme bets seem to be irrational.

⁸ Note, however, than in cases where the bet is proposed by a seemingly omnipotent bookmaker (say, an all-powerful genie), it can be argued that she could interfere with the truth of virtually any proposition the agent chose to consider.

pointed out above, if David has the possibility to stay clear from such a dangerous person, there are good reasons for him to do so – even if he remains certain about the truth of the proposition in question and his degree of credence does not change.

Perhaps one can imagine far-fetched science-fiction scenarios where an omnipotent but fair bookmaker routinely offers extreme bets to people, always respecting the terms of the bet. In this science-fiction world, being offered an extreme bet would be nothing surprising, and it seems that agents would have no reasons to distrust the bookmaker. It is not clear, however, that our actual-world intuitions can be reliably extrapolated to this type of far-fetched, distant world. It could well be that, if we lived in such a science-fiction scenario, we would regard many of the extreme bets offered as rationally acceptable. Be this as it may, the considerations in this section would still offer an explanation of why, in worlds similar enough to ours, it is rational to be reluctant to accept extreme bets on known propositions.

So, my suggestion is that the intuition that it is generally rational to reject extreme bets can be accounted for, at least to a large extent, by considering how the circumstances of the agent are altered by the fact that a bookmaker is willing to offer such a peculiar bet. I do not claim that all cases will be explainable in the exact same way. Perhaps there are further factors involved in other examples (say, some people may be reluctant to explicitly put certain valuable things at stake, even when they know that it is a mere pretence and there is no real risk). Nevertheless, I am confident that in most examples of extreme betting on known propositions where we find the relevant intuition of irrationality, it will be possible to come up with plausible stories similar to the one presented above. What I will show in the next section is that the apparent irrationality of extreme bets on known propositions vanishes in the absence of these sorts of distorting factors related to the presence of a bookmaker explicitly proposing the bet. In many ordinary situations, we are willing to act in ways that would have disastrous consequences if certain propositions were false. In other words, we are perfectly happy to take (implicit) extreme bets on the truth of a large number of mundane, empirical propositions.

3. RATIONAL EXTREME BETTING

I have argued that the reasons to avoid extreme bets on known propositions are given by features of the bookmaker proposing the bet: such a bookmaker is dangerous and untrustworthy; moreover the fact that she proposes the extreme bet could be evidence in

favour of freakish possibilities. If this is so, then when these features about the bookmaker are absent it should be rational to accept extreme bets on known facts. But, how can there be betting behaviour without someone proposing and enforcing the bet? A possibility is to see ordinary intentional agency as a form of implicit betting behaviour without a second party proposing the bet.⁹

I submit that rational extreme betting takes place, in an implicit way, in most of our ordinary actions. Many trivial actions would have terrible consequences if some of the success-conditions underlying them turned out not to obtain. For instance, going to the kitchen for a cup of tea will end badly if the kitchen floor collapses. Sometimes the agent is explicitly aware that her action's outcomes would be calamitous if certain conditions failed to take place. In these cases, her behaviour can be seen as a form of implicit extreme betting on the presence of such conditions. This form of extreme betting is, very plausibly, often rational – in particular, it is rational in relation to beliefs that do not seem to be bestowed with any special type of infallibility.

Consider the following example. David and his young (and imaginative) daughter Julia are about to go to the supermarket across the street. Julia has been reading about dangerous animals and is worried that a lion might eat them while crossing the street. David reassures her: there is no lion in the street, otherwise they would hear people screaming. Julia is satisfied with this and they safely cross the street to the supermarket.

David's behaviour is clearly rational. He is aware that if there happened to be a lion in the street, going to the supermarket would be very dangerous. However, David is confident enough that there are no lions around, so he rationally decides to go to the supermarket. David has good evidence to believe that there are no lions in the street: he has not heard or seen anything unusual in the street; lions do not live in this area; there are not even zoos in the surroundings. Moreover, nothing makes him suspect that, after all, perhaps there is some danger waiting in the streets. He has no reason, therefore, to

⁹ I will think of a decision as involving an implicit bet on p when the structure of winning rewards (if p is true) and losing costs (if p is false) associated with the decision is analogous to that of a certain bet on p , although there is no bookmaker explicitly proposing and enforcing such a bet. The idea that our ordinary actions can be regarded as implicit bets is already present in Ramsey ([1926]1978, p. 85): 'all our lives we are in sense betting. Whenever we go to the station we are betting that a train will really run.'

refrain from going to the supermarket. In this way, it is rational for David to behave as he would behave on the assumption that there is not a lion in the street (i.e. as he would behave if he assigned credence 1 to such a proposition). In other words, David is rational in implicitly accepting the extreme bet underlying this action – to wit, a bet on the truth of the proposition that there is not a lion in the street, with the winning outcome of getting his shopping done and the losing outcome of being attacked by a lion.

Now, things would be different if a magic genie or an omnipotent bookmaker with dubious intentions had explicitly offered David to bet his and Julia's lives on the truth of the belief that there is no lion in the street (with the winning outcome of getting his shopping done by the bookmaker). In such a case, David might have reasons to suspect that there is actually a lion in the street. If the bookmaker offers this extremely weird bet, perhaps it is because she has seen the lion, or has released it herself. Furthermore, David may be reasonably uncertain about the reaction of the almighty bookmaker if David wins: one does not want to find out how a murderous, unpredictable individual like this behaves when losing. In these conditions, it may well be rational for David to reject the bet. Thus, the same extreme bet can be rational when implicit in action, but irrational if explicitly proposed by an agent able to enforce it.

Another example. Back from the supermarket, David pours a glass of milk for Julia. If the liquid in the bottle were poisonous, it would be a very bad idea to give a glass of it to Julia. However, David knows that the bottle does not contain poison, but a perfectly drinkable, safe liquid: milk (it says so in the bottle, which he just bought in the supermarket). By offering Julia the glass, David is implicitly accepting a bet with abysmal losing costs: if he is wrong about the liquid not being poisonous, Julia's life will be in danger. Nevertheless, given what he knows, his behaviour seems rational. This is an (implicit) extreme bet that is rational to accept.

Actually, there seem to be many cases where people are ready to accept implicit extreme bets on the truth of p while explicitly acknowledging that there is some (small) probability that p is false. After all, people take planes and trains all the time, despite knowing that there is a small chance there will be an accident. Imagine, for instance, that David and Julia are taking the train to have a picnic by the lake (which they slightly prefer to a picnic by the park near their house). Julia has been reading about train disasters, and she is afraid that the bridge over the river will break when the train is

crossing it. However, David is familiar with train accident statistics, so he knows that, while it is possible that such an accident happens, the probability is extremely low. Thus, David assures Julia that the train is safe and they happily go and have a picnic by the lake. By deciding to take the train, David is implicitly accepting a bet with terrible losing costs (he and Julia would be involved in a catastrophic train accident) and with a comparatively very moderate winning reward (they only have a slight preference for picnics in the lake over picnics in the park). Nevertheless, David's behaviour seems rational. Again, this is an (implicit) extreme bet that is rational to accept, despite the fact that David assigns some (admittedly, practically negligible) probability to the losing outcome.

By contrast, it far less clear that the analogous explicit bet is rationally acceptable. Imagine that David is watching a train about to cross the bridge (he is not aboard), and a powerful bookmaker offers him to bet his and his daughter's lives on the truth of the proposition that the bridge will not break when the train crosses it, with the winning reward of having a picnic in the lake rather than in the park. It seems reasonable for David to stay clear of such an offer and such a bookmaker. This example shows that, even when the agent assigns some small probability to the losing outcome, there are implicit extreme bets that are rationally acceptable despite having the sorts of losing costs that seem to make analogous explicit bets unacceptable.¹⁰ Note, moreover, that we can stipulate that the probability initially attributed by David to the bridge breaking is the same in the explicit and the implicit versions of the bet (e.g. a probability matching the statistics of train accidents in that type of bridge). This suggests that the unacceptability of the explicit version of the bet cannot be explained just in terms of the stakes and the probabilities initially assigned to the losing outcome, since these are the same in the acceptable implicit bet.

Anyway, implicit extreme bets are not only acceptable when the costs remain below a certain threshold. We can set the losing costs as high as desired, and still find examples of rational implicit bets with those costs. Imagine that David is a historian,

¹⁰ According to standard decision-theory, when the agent assigns some probability to the losing outcome, the implicit bet will become unacceptable for sufficiently high losing costs (leaving winning rewards fixed). Perhaps in David's example we can find costs high enough to make it irrational to accept implicit bets about the safety of the bridge. Imagine that David is commissioned with planning the transport of a delicate nuclear device, which could destroy the whole city if there is an accident. Maybe in this case the rational thing for David to do is to avoid the bridge and send the convoy through a slightly longer, but slightly safer route.

investigating an old civilization that held the belief that eating beans for breakfast will condemn you and all your loved ones to earthly misfortune and eternal damnation. In front of his morning bowl of beans, David considers what it would be like to live in a world governed by such religious commands, and then he proceeds to eat his beans. Again, David is implicitly accepting an extreme bet that would be disastrous to lose (i.e. a bet on the truth of the old civilization's belief). He is rational in doing so because, knowing what he knows about religious practices and about this old civilization (their religious precepts were just written by common people, as part of their mythological literature), David has no reasons to think that he and his loved ones will be damned if he eats the beans.

What these examples show is that, for losing costs as high as we can imagine, most people are rationally disposed to take *implicit* extreme bets with such costs on the truth of a wide range of worldly, ordinary propositions. We are generally willing to act in ways that would be disastrous if those propositions were false, and that would only bring moderate benefits if true. It is worth stressing that this applies to common beliefs such as "There is milk in the glass" or "No car is approaching", and not just to a narrow set of special propositions (say, cogito propositions or logical truths). It is easy to find ordinary propositions that can figure in rationally acceptable extreme implicit bets (regardless of how extreme these bets are made). Thus, the scope of the alleged sceptic implications of extreme betting for views on which knowledge involves credence 1 would be severely reduced – pace authors like Williamson (2000, p. 213), who, as we saw above, claims that few propositions would pass the test posed by extreme betting. Moreover, it should be noted that the propositions involved in this form of extreme betting implicit in action do not enjoy a particularly strong epistemic support that sets them apart from other ordinary, worldly beliefs. It seems, therefore, that all those other ordinary beliefs with similar support are also good candidates for knowledge, even if perhaps they are not directly involved in extreme bets implicit in the agent's actions (think for instance of the belief that no car is approaching the crosswalk, when you are just looking from your window and not planning to cross the street).

One may think that the long-term costs of rejecting all these implicit extreme bets are very high (namely, practical paralysis), so that it is rational to adopt the general policy to accept such bets – even if an isolated instance of this type of bet would keep being irrational. Indeed, it may be argued that any decision involves an implicit extreme

bet on the truth of some propositions, since any option that one may choose (including inaction) will have terrible costs if certain conditions fail to obtain. So, while going to the kitchen for a glass of milk would be disastrous if the kitchen's floor collapses, not going to the kitchen and staying in your bedroom instead would also be catastrophic were your bedroom's ceiling to fall down. If we allow fanciful possibilities like the ones involved in David's breakfast example, this idea generalizes. After all, it could be that some evil demon will punish David if he *refrains* from eating his morning beans. Thus, it would seem that we cannot help taking implicit extreme bets, in any choice we make. One may think that the explanation of the rationality of ordinary actions involving implicit extreme bets is that such bets are inescapable (even if it would be unadvisable to take this type of bet if we could avoid it). According to this view, the options that is rational to choose in our ordinary decision-making are not particularly desirable themselves, given that they are associated with extreme bets, but they are less bad than the other options available, which also involve extreme bets. Call this possible explanation of the rationality of implicit extreme bets the *inescapability explanation*.

There are a few things that can be said in response to this possible explanation. First, an analogous line of argument could be applied to explicit extreme bets, leading to the conclusion that such explicit bets are also often acceptable (and thereby dissolving the challenge that explicit extreme bets seem to pose to views in which knowledge involves credence 1). Those persuaded by the inescapability explanation should think as well that accepting an explicit extreme bet will be rational in many cases, given that rejecting it would also mean taking (implicit) extreme bets (the inescapability explanation relies on the idea that all possible courses of action involve extreme bets). Since most authors agree that explicit extreme bets are generally unacceptable, this is an unattractive consequence of the inescapability explanation. Anyway, if it were the case that explicit extreme bets are after all often rational, then these bets would stop being problematic for the thesis that knowledge involves credence 1 (remember, the original problem for this thesis was the implausibility of its apparent prediction that explicit extreme bets on known propositions are rational). To be sure, advocates of the inescapability strategy can try to explain the general irrationality of explicit extreme bets (but not of implicit ones) by appealing to features of these bets absent in their implicit counterparts, that is features of the bets having to do with their being explicitly proposed. However, in principle such an explanation would be available as well for those who think that

knowledge requires credence 1 (in section 2 I have proposed a possible explanation of this type).

Thus, the inescapability explanation, if correct, would not really work against the thesis that knowledge involves credence 1, but would actually support it against the challenge posed by explicit extreme bets. At any rate, there are reasons to think that the inescapability explanation is misguided. It seems that if such an explanation were on the right track, we should be far more cautious than we actually are. Arguably, not all extreme bets involved in ordinary actions are equally safe. The circumstances in which the extreme bet would be lost are more far-fetched in some cases than in others. In this way, the possibility that an evil demon will punish David for eating beans seems more outlandish than the possibility that there is a lion in the street, or that the ceiling will collapse. Under the assumption that actions involving extreme bets can be rational only because such bets are inescapable, one would expect that, when possible, agents would tend to choose those options associated with safer extreme bets. An advisable strategy would be to go for the safest extreme bets available, other things equal. Yet this does not seem to be the case. Consider the action of eating. In general, by eating one takes the implicit bet that the food to be eaten is not poisoned. Now, it can be argued that a way of avoiding poison in your food is to cook it yourself, so that you can check what goes into it. Thus, eating in a restaurant would be less safe than cooking your own food, at least regarding the possibility of being poisoned. By cooking your food you are still taking implicit extreme bets, but these bets are safer than the ones associated with going to restaurants. Therefore, according to the inescapability explanation, it would be wise to prepare your own food whenever your preference for going to a restaurant is not significantly higher than your preference for cooking (that is, whenever you are initially indecisive between the two options). However, in situations where one is torn between eating out and staying at home, because the reasons in favour and against the two options roughly balance each other out, one does not typically consider the possibility of being poisoned in order to make a decision. This is just not the sort of factor we generally take into account when deliberating about what to do.

Furthermore, according to the inescapability explanation, we should experience great anxiety in each ordinary decision we make (at least after reflecting about the potential costs involved). However, this does not seem to happen. Assume that it is true that ordinary actions are only rational because the extreme bets they involve are

inescapable, although such bets would be unadvisable considered in isolation. Then, in each decision we would be facing a choice among bad options, that is among options that considered in isolation are undesirable and dreadful. Certainly, it is possible to find oneself in such unfortunate situations, where the rational thing to do is to choose the lesser bad. Yet when this happens, we still feel the anxiety brought by facing the undesirable course of action chosen. Perhaps it is the most rational thing we can do, but having to do it is still distressful. Think for instance of the prisoner who is forced to choose between the firing squad and joining the army, in the midst of a terrible war. Arguably, it is rational for the prisoner to choose war, but this does not mean that she will not feel anxiety and distress at such dim prospects. The fact that one is in a situation where a bad choice is unavoidable does not make such a choice less distressing. In other words, the potential costs of one option are not cancelled out by the potential costs of the available alternatives. I will still have to face the possible costs of whichever option I end up choosing, so my position is not made better or less disquieting just because the alternatives are worse.

Clearly, being forced to accept among explicit extreme bets, like the one Mary proposes to David, would be generally distressing. So, if accepting an implicit extreme bet were itself as undesirable as accepting an explicit one, it seems that we should feel anguish and anxiety any time we made a trivial choice (at least, when we realized that such a choice entails taking an implicit extreme bet). Ordinary decisions would involve choosing among options that are all undesirable, when assessed in isolation. No doubt, this would be distressing. However, it is clear that most of us feel quite relaxed when making ordinary decisions (e.g. when deciding whether to drink a glass of milk), even when we are aware that there are far-fetched possibilities that would make those actions disastrous. As I see it, these considerations provide convincing reasons against the view that implicit extreme bets are rationally acceptable only because they are associated with any possible course of action, and thereby they are inescapable. It seems that many ordinary courses of action are rationally eligible because they are themselves worth pursuing, even if they involve implicit extreme bets, and not just because the alternatives are worse.

4. CONCLUSIONS

I have argued that the apparent irrationality of extreme betting does not actually have as much epistemological relevance as it is often thought. The cases in which extreme bets on seemingly known propositions are in fact irrational can be explained away (by appeal to features of the bookmaker enforcing the bet), without undertaking weighty commitments in the epistemological debate about knowledge, certainty and scepticism. Furthermore, many ordinary actions can be seen as rational instances of (implicit) extreme bets, concerning ordinary empirical propositions. Even when the costs of being mistaken are very high, agents often behave as they would on the assumption that their beliefs are true. Thus, the claim that (implicit) extreme betting is always irrational is false.

Acknowledgements

I thank María Jiménez Buedo, Natasha Mullin, David Teira, several anonymous referees and an audience at UNED university for their feedback and comments. This work has been supported by the Spanish Government research project FFI2014-57258-P ("Normative inferences and interferences in scientific research") and the Portuguese Foundation for Science and Technology project UID/FIL/00183/2013 ("Values in Argumentative Discourse").

REFERENCES

- Brown, J. (2008). Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs*, 42(2), 167-189. doi:10.1111/j.1468-0068.2008.00677.x.
- Clarke, R. (2013). Belief is credence one (in context). *Philosopher's Imprint*, 13, 1-18. <http://hdl.handle.net/2027/spo.3521354.0013.011>.
- Dodd, D. 2017. Belief and certainty. *Synthese*, 194(11), 4597-4621. doi:10.1007/s11229-016-1163-4.
- Eriksson, L., & Rabinowicz, W. (2013). The interference problem for the betting interpretation of degrees of belief. *Synthese*, 190(5), 809-830. doi:10.1007/s11229-012-0187-7.

- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford: Oxford University Press.
- Greco, D. (2013). Probability and Prodigality. *Oxford Studies in Epistemology*, 4, 82-107.
- Greco, D. (2015). How I learned to stop worrying and love probability 1. *Philosophical Perspectives*, 29(1), 179-201. doi:10.1111/phpe.12059.
- Hacking, I. (1965). *Logic of statistical inference*. Cambridge: Cambridge University Press.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *The Journal of Philosophy*, 105(10), 571-590. doi:10.5840/jphil20081051022.
- Maher, P. (1993). *Betting on Theories*. Cambridge: Cambridge University Press.
- Ramsey, F.P. (1926). Truth and Probability. In D.H. Mellor (ed.) *Foundations: Essays in Philosophy, Logic, Mathematics and Economics*, London: Routledge (1978), 58–100.
- Reed, B. (2010). A defense of stable invariantism. *Noûs*, 44(2), 224-244. doi:10.1111/j.1468-0068.2010.00738.x.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.
- Weatherson, B. (2005). Can we do without pragmatic encroachment? *Philosophical Perspectives*, 19(1), 417-443. doi:10.1111/j.1520-8583.2005.00068.x.
- Weatherson, B. (2012). Knowledge, bets, and interests. In Jessica Brown and Mikkel Gerken (eds.), *Knowledge ascriptions*. Oxford: Oxford University Press, 75-103.
- Wedgwood, R. (2012). Outright belief. *Dialectica*, 66(3), 309-329. doi:10.1111/j.1746-8361.2012.01305.x.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- Williamson, T. (2005). Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge. *Philosophical Quarterly*, 55(219), 213-235. doi:10.1111/j.0031-8094.2005.00396.x.
- Worsnip, A. (2016). Belief, Credence, and the Preface Paradox. *Australasian Journal of Philosophy*, 94(3), 549-592. doi:10.1080/00048402.2015.1084343.

