



Sex crimes and misdemeanours

Campbell Brown¹ 

Published online: 18 February 2019
© The Author(s) 2019

Abstract How wrong is it to deceive a person into having sex with you? The common view seems to be that this depends on the nature of the deception. If it involves something very important, such as your identity, then the wrong done is very serious. But if it involves something more trivial, such as your natural hair colour, then the wrong seems less great. Tom Dougherty rejects this view. He argues that sexual deception is always seriously wrong. In this paper, I present a response to Dougherty’s argument. I propose an analysis of the wrongness in deception according to which acts of deception, in sexual relations and elsewhere, may differ in their degree of wrongness, and some may not be seriously wrong.

Keywords Deception · Consent · Deal-breakers · Sex

1 Introduction

Among the various means by which we influence the behaviour of others, one of the least reputable may be deception. By deceiving a person, we may induce her to do what she would not have done otherwise, had she known the truth. Deception seems generally wrong. But the matter might not be black and white. Some acts of deception seem less wrong than others; some, perhaps, are not wrong at all. Suppose you give your arachnophobic friend a false assurance that there are no spiders in the woods so that she will enjoy a pleasant walk that otherwise she would have been too afraid to take. (The spiders are so well hidden that your friend will never know they are there.) This seems at most a minor wrong. On the other hand, if you

✉ Campbell Brown
c.f.brown@lse.ac.uk

¹ London School of Economics, London, UK

misrepresent yourself as a bank manager to an elderly pensioner so as to defraud him of his life savings, then this is very wrong.

Deception also occurs in sexual relationships. People sometimes exaggerate, embellish, or flat-out lie to get other people to have sex with them. Here, too, common opinion seems to recognise a spectrum of wrongness. Contrast two cases:

- Case 1 Alice, a Democrat, dislikes Republicans so much that she would never knowingly have sex with one. She meets a man who professes not to be a Republican, though in fact he is. They enjoy a pleasurable evening together, culminating in sex.
- Case 2 Bert's wife has an identical twin sister who tricks him into having sex with her by impersonating his wife.

Assume, in both cases, the sex is enjoyable and neither Alice nor Bert will later come to regret the encounter if they do not discover the deception. The wrong done to Bert seems much worse than that done to Alice. At least this I take to be the common view, what most people would think.

Dougherty (2013) challenges this orthodoxy, which he calls the "Lenient Thesis". He argues that all instances of deceiving a person into sex are seriously wrong. He summarises his argument as follows:

1. Having sex with someone, while lacking her morally valid consent, is seriously wrong.
2. Deceiving another person into sex involves having sex with that person, while lacking her morally valid consent.
3. Therefore, deceiving someone into sex is seriously wrong. (Dougherty 2013, 720)

I shall defend the Lenient Thesis. I aim to show that Dougherty's argument is unsound. However, rather than objecting directly to one its premises, I shall pursue a more indirect strategy. I begin by proposing an independent argument for the orthodox view. This will involve developing an analysis of wrongness in deception. I shall then consider which of Dougherty's premises to reject in light of this analysis.

Dougherty explicitly rejects a harm-based view of the wrongness in sexual deception. Clearly, if sexual deception is wrong just insofar as it causes harm, then some acts of sexual deception are more wrong than others, since some cause more harm. Dougherty, however, rejects this view. He argues that some harmless acts of sexual deception (e.g., where the deception goes undiscovered) are nonetheless seriously wrong. The wrongness in sexual deception is better explained, he argues, by the absence of valid consent (Dougherty 2013, 724–727). I am more sympathetic to the harm-based view. However, I shall, for argument's sake, set that view aside here. I aim to show that, even if we explain the wrongness in sexual deception in terms of consent, we may distinguish different degrees of wrongness.

2 Deal-breakers and degrees

Dougherty focuses on cases in which a person agrees to sex as a result of being deceived about a “deal-breaker”. A deal-breaker, roughly, is a fact about a person’s prospective sexual partner such that the person would not agree to sex were she to learn this fact. Since Alice would not knowingly agree to sex with a Republican, the fact that a person is a Republican is a deal-breaker for her. My aim in this section is to argue that not all deal-breakers are equal. Some are stronger than others.

2.1 An analysis of deal-breaker strength

Here is my proposal in broad outline. A deal-breaker, in Dougherty’s words, is “a feature of the sexual encounter to which the other person’s will is opposed” (Dougherty 2013, 719). But opposition is a matter of degree: a person’s will may more strongly oppose some features than others. The strength of opposition may be revealed in situations of uncertainty or risk, where the feature is not known either to be present or absent. We may then see how great a risk of choosing an option with this feature the person is willing to take. The greater is this risk, the weaker is her opposition, and hence the weaker is her deal-breaker.

Consider, for example, deal-breakers in a different context. For a vegan, the fact that a meal contains animal products is a deal-breaker with respect to eating it. A vegan is someone who would not knowingly consume animal products. But some vegans are more committed than others. Imagine two vegans, Fiona and Michael. One day, they visit their friend, Hugh, who has prepared a meal for them. Being somewhat absent-minded, however, Hugh is not sure whether he made his pasta with vegan pesto. “I’m pretty sure it said ‘vegan’ on the jar,” he says with little conviction. Michael is not deterred: “That’s good enough for me,” he says as he tucks in. But Fiona is more cautious: “Well, if you’re not certain, then I’d rather not take the risk.” Assume that, were there no doubt as to the vegan-friendliness of the meal, Fiona and Michael would be equally eager to consume it. They are equally hungry, equally fond of Hugh’s cooking, and so on.

We may infer from their choices that Fiona is the more committed vegan. Her will opposes the consumption of animal products more strongly than Michael’s does. Notice, however, the difference between Fiona and Michael is not *whether* the presence of animal products is a deal-breaker. Michael, like Fiona, would not eat Hugh’s pasta if he knew it contained animal products.¹ The difference, rather, is that in situations such as this, where they do not know either way, Michael is prepared to *risk* eating animal products, but Fiona is not. It seems quite natural to say, therefore, that this is a *greater* deal-breaker for Fiona than for Michael.

Using the tools of decision theory, we may define a measure of deal-breaker strength along these lines. We can ask *how great* a risk of eating animal products a

¹ One might dispute whether this fact alone suffices to make Michael a genuine vegan. But that’s not the issue here. What is important is that the presence of animal products *is* a deal-breaker for Michael, according to Dougherty’s definition.

person will take. The greater is this risk, the smaller is the deal breaker. The decision in the above example can be represented as follows:

Actions	States	
	Non-vegan (p)	Vegan ($1 - p$)
Eat	Satisfied & Immoral	Satisfied & Moral
Don't eat	Unsatisfied & Moral	Unsatisfied & Moral

There are two possible states of the world: either the food is *non-vegan* (i.e., containing animal products), or it is *vegan*. Let the probability of the former be p , and the latter $(1 - p)$. The agent may choose either to eat the food or not. The outcome of each action in each state is shown in the table. If the agent chooses to eat the food and it is non-vegan, then her hunger will be *satisfied*, but she will have acted *immorally*, by the ethical standards of veganism. (Assume she is a vegan on ethical grounds.) If she chooses to eat the food and it is vegan, then she will be both satisfied and moral. Finally, if she chooses not to eat the food, then she will be unsatisfied but moral in either case. (Assume there is no other food to eat.)

Let us write S for *satisfied* and M for *moral*. The presence of animal products is a deal-breaker for the agent if she prefers $S \& M$ to $\neg S \& M$, and prefers $\neg S \& M$ to $S \& \neg M$. If she is certain the food is non-vegan ($p = 1$), then she prefers not to eat it; but if she is certain the food is vegan ($p = 0$), her preference is reversed.

We may then consider intermediate cases, where $0 < p < 1$. These are cases of uncertainty or risk. We may assume that increasing the value of p will reduce the agent's inclination to eat the food. That is, if she prefers to eat the food for some value of p , then she will also prefer this for any lower value. We may then ask what is the maximum value of p for which she prefers to eat the food (or, more precisely, what is the least upper bound of the values of p for which this holds). The greater is this maximum value, the smaller is the deal-breaker. Fiona's maximum, for example, is lower than Michael's.

Suppose the agent's preferences are represented by an expectational utility function U . Then the strength of her deal-breaker is given by the ratio:

$$\frac{U(\neg S \& M) - U(S \& \neg M)}{U(S \& M) - U(S \& \neg M)}$$

The greater is this ratio, the smaller is the risk she is willing to take, and hence the stronger is her deal-breaker. To illustrate, suppose Fiona's and Michael's utilities are as shown in the following table:

	Utility			Deal-breaker strength
	$S \& M$	$\neg S \& M$	$S \& \neg M$	
Fiona	12	10	2	0.8
Michael	16	4	1	0.2

Fiona and Michael do not differ with respect to the *order* of the three outcomes: both say $S\&M$ is best, $S\&\neg M$ is worst, and $\neg S\&M$ is in the middle. The presence of animal products is thus at least a weak deal-breaker for both. Their wills oppose the consumption of animal products to at least some degree. However, they differ with respect to the location of the middle option vis-à-vis the other two: Fiona has it closer to the top; Michael has it closer to the bottom.² (A more general definition of this measure is given in the “Appendix”.)

This provides a precise numerical measure of deal-breaker strength on a ratio scale. It enables us to say, for example, that one deal-breaker is twice as strong as another. However, the measure rests on assumptions that may be controversial. It assumes, in particular, that a person’s “will” is aptly represented by her preferences over “lotteries”. In the above example, the agents’ preferences over lotteries involving some risk of consuming animal products are taken to show how much veganism *matters* to these agents. One reason to doubt this is that it is insensitive to the agents’ general attitudes to risk. If Michael is a thrill-seeker who likes to live life on the edge, whereas Fiona is more conservative, then Michael’s greater preparedness to risk eating animal products might not show that veganism matters less to him.

I am unsure what to say about risk aversion in this context. It does not seem obviously mistaken to say that risk-lovers generally have fewer or weaker deal-breakers than others. But if this is problematic, then the measure might be modified to avoid this implication. When we compare two agents with the same level of general risk aversion (or we are comparing a single agent in two different situations), the measure would remain as above. But in other cases, we might apply a variable “discount rate” to the agent’s utilities which would correct for differences in general risk aversion.

Still, even if such a modification could deal adequately with risk aversion, some may object, more broadly, that the decision-theoretic approach taken here is misguided. A person’s will is better understood in terms of the ordinary notion of *desire*, rather than the technical notion of preference employed by decision theorists and economists. When we say, for example, that veganism matters more to Fiona, we simply mean that she has a stronger desire to avoid eating animal products. This is easily comprehended without the need for utility functions and the like.³ This objection raises important issues about the psychology of desires and preferences which are beyond the scope of this essay. I should emphasise, however, that the *measure* of deal-breaker strength I have proposed is separable from the more general *analysis* of such strength on which it is based. This analysis says merely that, as a deal-breaker is a feature to which a person’s will is opposed, the strength of a deal-breaker correspondingly varies with the strength of the will’s opposition. Precisely how one elaborates this notion of a person’s will opposing a feature is a further issue. I have suggested one way of doing this, in decision-theoretic terms. A

² This does of course presuppose interpersonal comparisons of utility (in particular, cardinal comparisons of utility intervals).

³ Thanks to an anonymous referee for pressing me on this issue.

nice feature of this approach is that provides a way to define a measure of deal-breaker strength. But this is inessential for my argument. The essential point is merely that deal-breakers vary in strength.

A final note of clarification. On the analysis I am offering, certain patterns of preferences are what *determine* the strength of a deal-breaker; they are not mere *evidence* that the deal-breaker has a certain strength. Nonetheless, there *is* an important epistemic element in the overall account. We cannot directly observe a person's preferences, and therefore, on this account, we cannot directly observe her deal-breakers. Rather, we infer a person's preferences, at least in part, from her observed choice behaviour. When Fiona chooses not to eat the pasta, this reveals to us something about her preferences, and thereby something about her deal-breakers. I return to this issue below when discussing the question of wrongness in our two cases of deceiving into sex.

2.2 An alternative analysis

Before continuing, I would like to compare the analysis proposed above with a similar one suggested by Neil Manson, which also distinguishes stronger and weaker deal-breakers. The extent to which these analyses are incompatible is not entirely clear. But in any case, as I argue below, my proposal seems preferable, because it offers a deeper explanation.

Manson explains his analysis as follows:

The gradable voluntariness theory takes into account the fact that some kinds of deception may merely tip the balance in favor of deciding to consent. Where deception tips the balance, the relevant counterfactual—if R were to know that p she would not consent—is true, but the counterfactual is only made true by a relatively small set of very close possible worlds. Suppose R wouldn't want to have sex with older men. S , aware of this, dyes his hair to appear younger. Suppose R finds S attractive, witty, sexy. R willingly has sex with S . There are some close possible worlds where R knows the facts about S 's age and does not consent. But there are also close possible worlds where R knows the facts and does consent (she might, e.g., revise her extant preferences in light of new information: “Who knew? Not all old guys are monsters!”).

But deception can also lead the consentor to make a never-in-a-million-years type of decision. Suppose S deceives a devoutly religious R —who would rather die than have sex with a person outside her religion—into having sex, by lying about his religion. Even if R has other reasons for consenting to sex with S , there are no close possible worlds where R knows the truth about S 's religion and R consents. In this example, unlike the one above, we have a fantastically strong deal breaker. (Manson 2016, 419)

Suppose some fact F is a deal-breaker for R . To evaluate its strength, on Manson's account, we move from the actual world to the “closest” possible world where F is not a deal-breaker for R . The greater the distance we must travel, through modal space, the stronger is the deal-breaker. This raises two questions. First, when

surveying other possible worlds, which facts must be held fixed and which are allowed to vary? Second, how do we measure similarity, or “closeness”, between worlds? Though Manson does not explicitly supply these details, his parenthetical comment—“she might, e.g., revise her extant preferences”—suggests the following. The facts which are held fixed are those about the objects of R 's preferences (e.g., facts about S , in Manson's example), and the facts which are allowed to vary are those about R 's preferences. The relevant measure of similarity, then, is between preferences. The question is: holding fixed the other facts, how great a change in R 's preferences would be required for F to cease being a deal-breaker for R ?⁴

This raises the question of how to measure similarity between preference relations. Perhaps the simplest approach is to count the number of common pairs. For example, let \succsim and \succsim^* be preference relations on a set of options X . Then we may consider the cardinality of their intersection, i.e., of the set $\{\langle x, y \rangle \in X^2 : x \succsim y \wedge x \succsim^* y\}$. The larger this set, the greater the similarity between the preference relations. It is unclear, however, whether this will deliver the result Manson wants in his two scenarios, i.e., that there is greater preference change in the “age” case than in the “religion” case. In either case, merely reversing R 's preference over a single pair of options (i.e., having sex with S , and not doing so) may be sufficient for the relevant fact to cease being a deal-breaker. So this simple cardinality measure seems inadequate for Manson's purposes.

Rather, Manson seems to assume a measure of preference change that takes into account similarity between the *objects* of preferences. Such a measure may be fairly intuitive. For example, suppose I prefer beer to wine, and, among beers, I prefer ale to lager. Now consider two possible changes in my preferences. In the first, I continue to prefer beer to wine, but I reverse my preference regarding ale and lager. In the second, I continue to prefer ale to lager, but I reverse my preference regarding beer and wine. It seems natural to say that the latter change is greater. A plausible explanation is that ale is more similar to lager than beer is to wine.

In Manson's first scenario, the relevant preference change involves R 's preferences regarding older and younger men, whereas in the second scenario, it involves R 's preferences regarding, say, Christians and non-Christians. It might be said that, because there is in general more similarity between older men and younger men than there is between Christians and non-Christians, the latter preference change is greater, as required by Manson's analysis.

However, it is doubtful that this analysis will work in cases like that of Fiona and Michael above. In this case, the relevant preference change seems the same for Fiona and Michael. For both of them, the preference change involves their preferences regarding vegan food and non-vegan food. Whatever the degree of similarity between vegan food and non-vegan food, this will be the same for Fiona

⁴ It might be objected that this interpretation puts too much emphasis on a mere parenthetical remark. Perhaps, this remark notwithstanding, Manson does not intend closeness between worlds to be measured in terms of similarity between preferences, or not solely in such terms. I do not deny that there may be a more attractive interpretation. Notice, however, the concept of a deal-breaker suggests that preferences must at least be a central part of the story. Dougherty defines a deal-breaker in terms of the agent's *will*, and it is natural to think of this as closely related to, if not constituted by, the agent's preferences.

and Michael. Yet, as I said above, it is plausible that Fiona's deal-breaker is stronger. It is unclear, therefore, whether Manson's analysis can give a satisfactory explanation of this case.

In any event, I think there is a good reason to prefer an analysis like the one I have proposed to Manson's alternative. Suppose, for argument's sake, that there *is* a difference in the "size" of preference change required for the presence of non-vegan ingredients to cease being a deal-breaker for Fiona and Michael. Presumably, if this is so, then it is due to some underlying difference in Fiona's and Michael's *actual* preferences. A deeper explanation, therefore, will be one which identifies what this actual difference may be. This is what my proposal attempts to do. It explains the difference in deal-breaker strength between Fiona and Michael in terms of their actual preferences. Now, it might be the case that, given this difference in their actual preferences, the size of preference change required for Fiona is greater than that required for Michael, in which case Manson's analysis may agree with mine. But then mine would seem to give the deeper explanation. On the other hand, if the two analyses do not agree, then mine seems more plausible.

2.3 Local versus global deal-breakers

The analysis above concerns what might be called "local" deal-breakers. It concerns the extent to which a property is a deal-breaker in particular circumstances. This might vary in different circumstances. Fiona might be more willing to risk eating Hugh's pasta were she very hungry; Michael might be less willing to do so were Hugh known to be a bad cook. This might not sit well with our ordinary notion of a "deal-breaker". We normally think of deal-breakers as being more robust, not so sensitive to such changes in circumstances. It would be absurd to say, for example, "I'm a vegan, but only when I'm not hungry."

To better capture the ordinary notion, we may also consider a "global" dimension of deal-breaking strength. A property might be a (local) deal-breaker in one situation but not in another. We may thus think of its global strength as determined by how broad or narrow is the range of situations in which it is a deal-breaker. A maximally strong global deal-breaker—call it a "fundamental" deal-breaker—is a local deal-breaker in *all* possible situations. Presumably, the presence of animal products is not a *fundamental* deal-breaker even for most vegans. Few, I assume, would starve to death rather than compromise their vegan scruples, for example. We may be sceptical about the existence of *any* fundamental deal-breakers. We may believe that "everyone has a price". For example, most people say they would never eat human flesh, yet there are well-known cases of ordinary people resorting to cannibalism in extreme circumstances.

Among non-fundamental global deal-breakers, we may distinguish weaker and stronger: X is globally at least as strong as Y if and only if every situation in which Y is a deal-breaker is also a situation in which X is a deal-breaker, i.e., $D_X \supseteq D_Y$, where D_X and D_Y are the sets of situations in which, respectively, X and Y are deal-breakers. For example, the presence of meat may be for many vegans a stronger deal-breaker than the presence of animal products. If a person would not choose to eat animal products in a given situation, then *a fortiori* she would not choose to eat

meat in the same situation. However, there may be situations in which, though she would not choose to eat meat, she would choose to eat, say, cheese. Notice, this gives us only a *partial* ordering of deal-breakers. In some cases, it may be that neither $D_X \supseteq D_Y$ nor $D_X \subseteq D_Y$ holds. In such cases, we may want to compare the relative size of D_X and D_Y , but this will be difficult if, as seems likely, D_X and D_Y are both infinite.

As I say, this global definition may be closer to our ordinary notion. In what follows, however, I will focus solely on local deal-breakers, since this seems closer to what Dougherty has in mind. Moreover, it may be argued that, from the standpoint of moral evaluation, local deal-breakers are more relevant, since the moral evaluation of an action should attend to the particular features of the context in which it occurs.

3 Wrongness in deception

What should we say about the wrongness in deception when we recognise that deal-breakers vary in strength? The following seems a plausible view. If a person acts in a way that morally requires the consent of another person, and this consent is secured through deception regarding a deal-breaker, then the wrongness of this action correlates positively with the strength of the deal-breaker. Actions which are enabled by deception with respect to strong deal-breakers are more wrong those enabled by deception with respect to weak deal-breakers. Moreover, some deal-breakers are so weak that the associated wrong is not serious. I begin by giving a general account of degrees of wrongness in deception, before applying this to the central case of deceiving people into sex.

3.1 Degrees of wrongness

Suppose that Hugh knowingly serves pasta containing non-vegan pesto to his vegan guests. When they express doubts about the vegan-friendliness of the food, he reassures them by displaying an empty jar of vegan pesto which he claims to have used in preparing the meal (though in fact he merely emptied the contents into the rubbish bin). Then, it seems, Hugh has acted wrongly towards both Fiona and Michael. But, on my proposal, the wrong done to Fiona is greater, because her deal-breaker is stronger. This seems the right verdict. Hugh has, through his deception, induced Fiona and Michael to act against their own preferences: they have eaten animal products when they would have preferred not to do so. However, Fiona has acted against her preferences to a greater extent than Michael has acted against his. The avoidance of eating animal products *matters more* to Fiona. This is why Michael is willing to run a greater risk of eating animal products. Fiona's preference in this case is *stronger* than Michael's. Since Hugh has induced Fiona to act against a stronger preference, he has wronged her more. In Dougherty's terms, Hugh has engineered an outcome to which Fiona's and Micheal's

wills are opposed. But, since Fiona's will is more strongly opposed to this outcome, the wrong done to her is greater.

Now suppose that Michael is a *very* half-hearted vegan. He routinely orders meals at restaurants after making only minimal efforts to confirm they are vegan-friendly. "It's probably okay", he says, though the information provided on the menu is inconclusive. Although he would not knowingly consume animal products, he has few qualms about risking it when he doesn't know either way. His deal-breaker is *very* weak. And suppose Hugh knows this, having observed Michael's behaviour in restaurants. Perhaps it is still wrong for Hugh to deceive Michael into eating non-vegan food, but, I would say, it is not clearly a serious wrong.

This picture fits nicely with the general analysis of degrees of wrongness I have developed elsewhere (Brown 2016).⁵ On this analysis, the degree of wrongness of an action depends on how similar it is to a right (non-wrong) action. How different would the action need to be for it to be right? For a given action A , consider an action A^* such that (i) A^* is right, and (ii) A^* is at least as similar to A as any right action.⁶ The more similar A^* is to A , the less wrong is A .⁷ Suppose, for example, A is an action of wrongful deception about a deal-breaker. Then A^* may be an act that is as similar as possible to A except it involves no deception. The relevant dimension of similarity here may be the amount, or level, of deception involved in the action. It seems natural to say that the amount of deception involved in A depends on the strength of the deal-breaker. The stronger this is, the more deception is involved, and hence the more dissimilar A is from A^* , and the more wrong is A .

It should be emphasised that, on this analysis, the wrongness of acts of deception need not be correlated with the amount of harm they cause. Rather, wrongness depends on the strength of deal-breaker involved. Of course, deal-breaker strength does depend on the agent's preferences. So, if harm simply consists in the frustration of preferences, then it may be that deception with respect to stronger deal-breakers causes greater harm. However, this is a non-essential feature of the view. Given a non-preference-based account of harm—say, a hedonistic one—it may be that, in some cases, deception with respect to a stronger deal-breaker causes less harm than deception with respect to a weak deal-breaker.

3.2 Deception and sex

Return now to our examples of sexual deception. What does my proposal say about these?

Consider first Case 1, in which Alice is deceived into having sex with a Republican. We may plausibly infer that Republicanism is only a weak deal-breaker for Alice. Everybody knows that in negotiations leading to casual sex, people are

⁵ Strictly speaking, this analysis is of degrees of *rightness*. However, applying it to wrongness instead is straightforward.

⁶ This involves a slight simplification, since there might not exist an action satisfying these two conditions. I suggest a more complicated version of the analysis to accommodate this possibility.

⁷ If A is itself right, then we may have $A^* = A$. In that case, A is not wrong at all, because there is zero dissimilarity between a thing and itself.

not always entirely candid. A certain lack of candour might even be considered a shared background assumption in this context, similar to that assumed in, say, poker games or reference letters. When your poker opponent grimaces, apparently revealing another bad hand, you know not to take this at face value. When a job candidate's reference letter declares that she "has the greatest philosophical mind since Aristotle", you know to interpret this a meaning something closer to "she has read some books by Aristotle." Similarly, when a guy you met at a bar tells you he's 36, you know he's probably in his forties. If Alice has a one-night stand with a man—even one who professes not to be a Republican—then she willingly takes a nontrivial risk of having sex with a Republican. That he claims not to be a Republican is, in this context, only weak evidence that he is not one. That she is nonetheless willing to take this risk shows that Republicanism is not a strong deal-breaker for her. If she is wronged, it is, therefore, not a serious wrong.

I should stress that this conclusion is defeasible. On the basis of observing Alice's choice behaviour, and certain background assumptions about her credences, we may make reasonable inferences about her preferences over relevant "lotteries", or uncertain prospects. These preferences, on the analysis proposed here, determine the strength of her deal-breaker. It is possible, however, that further information would show her that her preferences are different, and therefore that her deal-breaker is, say, stronger.⁸ A related question is what Alice's deceiver (the disguised Republican) may reasonably infer about the strength of her deal-breaker. If he reasonably infers that it is weak when in fact it is strong, then we might say that his deception is less blameworthy, though still seriously wrong.⁹

Consider now Case 2, in which Bert is tricked into having sex with his wife's identical twin. Cases like this are extremely rare—certainly much rarer than cases like Alice's, involving run-of-the-mill deceptions. It is entirely reasonable for Bert to believe, with high confidence, that the woman with whom he agrees to have sex is his wife. He does not willingly take any significant risk of having sex with a different person. Moreover, we may assume that, were he to begin to suspect that it was not his wife, he would end the encounter immediately. That this woman is not in fact his wife is therefore a strong deal-breaker for him, and the wrong done to him is correspondingly a serious one.

3.3 Core and periphery

Dougherty argues that, in order to accept the Lenient Thesis, we must distinguish "core" and "peripheral" features of a sexual encounter (Dougherty 2013, 729).

⁸ Consider a version of this case in which Alice goes to very considerable lengths to confirm that her prospective sexual partner is not a Republican: she interviews his family and friends; she calls the local branch of the Republican party to ask about its membership list; and so on. And suppose that he goes to equally great lengths to conceal his Republicanism: he instructs his family and friends to lie to Alice, and so on. In this case, Alice's deal-breaker is much stronger, and the wrong done to her correspondingly more serious.

⁹ That deceiving into sex may sometimes be blameless yet still seriously wrong is discussed by Dougherty (2013, 739–743).

Serious wrongdoing occurs in cases of sexual deception involving core deal-breakers, but not in those involving peripheral deal-breakers. He considers two approaches to drawing this distinction. The first he calls “objective,” because it is insensitive to individuals’ own views about what is important. For example, identity may be deemed a core feature, and political affiliation a peripheral feature, even though for some individuals, identity may be less important than political affiliation. The problem with this approach, he argues, is that it entails an unacceptably “moralistic conception of sex” (Dougherty 2013, 730). It imposes certain views of sexual morality on people who do not subscribe to these views. Dougherty favours instead a “subjective” approach. He writes: “The most principled way to [distinguish core and peripheral features] is to distinguish the features that someone considered relevant to her decision to have sex from those that she considered irrelevant” (Dougherty 2013, 731). He then points out that any deal-breaker must, by definition, concern a feature that the individual considers relevant. Thus *all* deal-breakers are core features.

Notice, however, that the analysis I have proposed is, by Dougherty’s lights, *subjective*. What is a strong deal-breaker for one person may be only a weak deal-breaker, or no deal-breaker at all, for another. Even if core features are distinguished from peripheral features by what individuals consider relevant, it does not follow that this is a simple binary matter. For relevance may itself be a matter of degree. When I buy socks, their colour is relevant to my decision, but not *as relevant* as their size. When Alice chooses with whom to have sex, presumably many features are relevant to her decision. One relevant feature is whether her prospective sexual partner is a Republican. But it cannot be very high on her list of relevant features, since she is willing to have sex with a man without good evidence that he is not a Republican. It is relevant, but not very relevant. The class of core features, which give rise to serious wrongdoing, may thus be restricted to a proper subset of the relevant features, including only those with sufficiently high relevance for the person deceived.

3.4 Consent

I have argued above that sexual deception is not always seriously wrong. Where does this leave the issue of consent? Recall Dougherty’s argument:

1. Having sex with someone, while lacking her morally valid consent, is seriously wrong.
2. Deceiving another person into sex involves having sex with that person, while lacking her morally valid consent.
3. Therefore, deceiving someone into sex is seriously wrong.

Which of the premises do I reject? This is not an easy question. It depends both on the relation between deception and consent, and the relation between consent and wrongness. On Dougherty’s view, deception always invalidates consent, and sex without consent is always seriously wrong. One might reject either of these claims.

First, one might say deception invalidates consent only when it involves a sufficiently strong deal-breaker.¹⁰ This would be to reject Premise 2. However, this may seem to misunderstand the widely accepted information requirement for valid consent. The reason why valid consent entails *informed* consent is that a person cannot validly consent to something if she does not know what she is consenting to. A person cannot, for example, validly consent to participate in risky drug trials if she does not know the risks. But if a person is deceived about a deal-breaker, even a very small one, then she does *not* know what she is consenting to. Alice, for example, does not know that she is consenting to have sex with a Republican. It seems difficult to maintain, therefore, that her consent is nonetheless valid.

Second, one might say non-consensual sex is seriously wrong only when it results from deception involving a sufficiently strong deal-breaker. This would be to reject Premise 1. However, one might reasonably be reluctant to say that non-consensual sex is sometimes only a minor wrong. As Dougherty point out, non-consensual sex is taken by some to be the definition of rape (Dougherty 2013, 721), yet surely rape is *never* only a minor wrong.

There is, however, a third possibility. We might say that violations of consent also come in degrees. The wrongness of non-consensual sex depends on *how* non-consensual it is. Sex that results from deception involving a strong deal-breaker is non-consensual to a high degree, and is therefore seriously wrong, whereas sex that results from deception involving a weak deal breaker is non-consensual to a low degree, and is therefore not seriously wrong. On this approach, we should say that Dougherty's argument is equivocal. What is meant by "having sex with someone while lacking her morally valid consent"? Does this refer to sex that is non-consensual to a high degree (involving a strong deal-breaker), or to a low degree (involving a weak deal-breaker). If the former, then we should reject Premise 2; if the latter, Premise 1.

The idea that acts may be non-consensual to a greater or lesser degree seems fairly intuitive. Suppose I employ you to paint my house. You show me the tins of paint you will use, which, according to their labels, are the colour I have carefully selected, a particular shade of white called "Snow White." If you have surreptitiously filled the tins with cheaper paint of a slightly different colour, "Periwinkle White," then you act to some extent without my consent when you paint my house. But not to the *same* extent as if the paint was bright purple. It seems natural to say that, though both acts of house painting are to some extent non-consensual, the latter is more so than the former. It also seems natural to say that latter is more wrong than the first.

I should stress that this argument is intended to apply only to cases in which consent is invalidated by deception. It does not apply to cases where a person is forced or coerced to do something against her will. Deception and force seem importantly different. In cases of deception, we may say that the victim *thought* she was consenting to one thing, X, when in fact she consented to something else, Y. In considering the extent of non-consensual action, it therefore makes sense to take

¹⁰ This I take to be the position of Manson (2016).

into account how great the difference is between X and Y . In a case of force, however, the victim may know perfectly well what she has and has not consented to. For this reason, the account of degrees I have suggested seems less appropriate in cases of force.

4 Conclusion

I have argued that acts of wrongful deception, in sexual relations and elsewhere, may vary in their degree of wrongness. In some cases, the degree of wrongness might be so small that we would not regard the deception as “seriously” wrong. My argument rests on a proposed analysis of this variation, whereby the degree of wrongness depends on the strength of the deal-breaker involved, as may be revealed by a person’s behaviour in risky situations. This view can, I have argued, be reconciled with a plausible account of consent. None of this should be taken as a general defence of deception. Many acts of deception surely *are* seriously wrong, and my analysis is entirely consistent with this. It is also compatible with the view that all acts of deception are wrong to at least some degree.

Acknowledgements I am grateful to Tom Daugherty and Hugh Lazenby for comments on earlier drafts of this paper, and also to an anonymous referee for this journal.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix

This appendix gives a more general treatment of the analyses suggested in Sects. 2.1 and 2.3.

Local deal-breakers

In my formal analysis, deal-breakers will be propositions. Let Ω be a finite set of states (finite for simplicity). A proposition X is a subset of Ω . Let A be a set of actions. For each $a \in A$, let there be a value function $V_a : \Omega \rightarrow \mathcal{R}$, where $V_a(\omega)$ represents the value of performing a in ω .

The agent’s *ex ante* degrees of belief, prior to receiving any testimony, are given by a regular probability function P on 2^Ω . (By saying P is regular, I mean that $P(X) = 0$ only if $X = \emptyset$.) The *ex ante* expected value of action a is thus

$$\mathbb{E}_P[V_a] = \sum_{\omega \in \Omega} V_a(\omega)P(\omega) \quad (1)$$

Now suppose the agent receives testimony that proposition X is true (where $X \neq \emptyset$), and updates her degrees of belief by conditionalising on X . So her new probability function is P_X , where, for all Y ,

$$P_X(Y) = P(Y|X) \quad (2)$$

The new expected value of a is thus

$$\mathbb{E}_{P_X}(V_a) = \sum_{\omega \in X} \frac{V_a(\omega)P(\omega)}{P(X)} \quad (3)$$

We may contrast this with the case where the agent instead receives testimony that X is false, and so conditionalises on \bar{X} , the negation (or complement) of X . It may be that the agent's preference regarding actions a and b "reverses" when switching from the former case to the latter: when she conditionalises on X , she prefers a to b , whereas when she conditionalises on \bar{X} , she prefers b to a . This means that X is a deal-breaker for her.

More precisely, X is a deal-breaker favouring a over b if the following two inequalities hold

$$\mathbb{E}_{P_X}[V_a - V_b] > 0 \quad (4)$$

$$\mathbb{E}_{P_{\bar{X}}}[V_a - V_b] < 0 \quad (5)$$

An immediate consequence of these definitions is that if X is a deal-breaker favouring a over b , then \bar{X} is a deal-breaker favouring b over a . And the strength of the former is inversely related to the strength of the latter.

Suppose that X is in fact false, but the agent is deceived into believing it is true. Then she will prefer a to b , but would have had the reverse preference had she known the truth. In this way, deception with respect to a deal-breaker may induce a person to act against her "true" preferences.

The above analysis is somewhat crude. It assumes testimony has the effect of raising the agent's credence in some proposition to the level of *certainty* (e.g., $P_X(X) = 1$). But of course this is unrealistic. So suppose now that the agent receives testimony that is insufficient to make her certain that X true, but rather leads her to assign X probability λ . If she then updates her degrees of belief according to Jeffrey Conditionalisation, her new probability function will be $P_{X,\lambda}$, where for all Y ,

$$P_{X,\lambda}(Y) = \lambda P(Y|X) + (1 - \lambda)P(Y|\bar{X}) \quad (6)$$

And the expected value of a becomes

$$\mathbb{E}_{P_{X,\lambda}}[V_a] = \sum_{\omega \in \Omega} V_a(\omega)(\lambda P(\omega|X) + (1 - \lambda)P(\omega|\bar{X})) \quad (7)$$

$$= \lambda \mathbb{E}_{P_X}[V_a] + (1 - \lambda) \mathbb{E}_{P_{\bar{X}}}[V_a] \quad (8)$$

We may consider the difference $\mathbb{E}_{X,\lambda}[V_a - V_b]$ for varying values of λ . This difference is given by

$$\begin{aligned} \mathbb{E}_{X,\lambda}[V_a - V_b] &= (\lambda\mathbb{E}_{P_X}[V_a] + (1 - \lambda)\mathbb{E}_{P_{\bar{X}}}[V_a]) \\ &\quad - (\lambda\mathbb{E}_{P_X}[V_b] + (1 - \lambda)\mathbb{E}_{P_{\bar{X}}}[V_b]) \end{aligned} \tag{9}$$

$$= \lambda\mathbb{E}_{P_X}[V_a - V_b] + (1 - \lambda)\mathbb{E}_{P_{\bar{X}}}[V_a - V_b] \tag{10}$$

Suppose X is a deal-breaker favouring a over b , as defined above. So (4) and (5) both hold. It follows that

$$\text{sign}(\mathbb{E}_{X,\lambda}[V_a - V_b]) = \begin{cases} 1 & \lambda > r \\ 0 & \lambda = r \\ -1 & \lambda < r \end{cases} \tag{11}$$

where

$$r = \frac{\mathbb{E}_{P_{\bar{X}}}[V_b - V_a]}{\mathbb{E}_{P_{\bar{X}}}[V_b - V_a] + \mathbb{E}_{P_X}[V_a - V_b]} \tag{12}$$

Thus the ratio r represents the lowest probability such that the agent may assign this probability to X without preferring b to a . We may therefore use r as a measure of the strength of X as a deal-breaker favouring a over b . The *smaller* is r , the *stronger* is the deal-breaker.

Global deal-breakers

Clearly, on the above analysis, whether a proposition is a deal-breaker depends on what I called the *ex ante* probability function. X may be a deal-breaker relative to one probability function P , but not relative to another P' . In this way, the analysis concerns what I earlier called *local* deal-breakers. These deal-breakers are localised to particular circumstances, represented here by particular *ex ante* probability functions.

As I suggested, we can also consider *global* deal-breakers. The deal-breaking strength of a proposition, in a global sense, depends on how broad or narrow is the class of circumstances relative to which it is a deal-breaker in the local sense. At one extreme, we have what I called *fundamental* deal-breakers. X is a fundamental deal-breaker if, for every (regular) probability function P , X is a deal-breaker relative to P . For actions a, b , define

$$\Omega_{a > b} = \{\omega : V_a(\omega) > V_b(\omega)\} \tag{13}$$

Thus $\Omega_{a > b}$ is the set of all states in which performing a is better than performing b . Then X is a fundamental deal-breaker, favouring a over b , if and only if $\Omega_{a > b} = X$ and $\Omega_{b > a} \neq \emptyset$.

We can also give a partial order of propositions representing their strength as global deal-breakers. The global strength of X as a deal-breaker favouring a over b is at least as great as that of Y if, for every probability function P , if Y is a deal-breaker favouring a over b relative to P , then so is X . This will be the case if $X \cap \Omega_{a > b} \supseteq Y \cap \Omega_{a > b}$ and $X \cap \Omega_{b > a} \subseteq Y \cap \Omega_{b > a}$.

References

- Brown, C. (2016). The Rightest Theory of Degrees of Rightness. *Ethical Theory and Moral Practice*, 19(1), 21–29.
- Dougherty, T. (2013). Sex, lies, and consent. *Ethics*, 123(4), 717–744.
- Manson, N. C. (2016). How not to think about the ethics of deceiving into sex. *Ethics*, 127(2), 415–429.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.