

Deception and Consent

Tom Dougherty

Draft. Forthcoming in the Routledge Handbook of the Ethics of Consent

Introduction

Even though April Fool's day is an inauspicious day for a wedding, you still would not expect the marriage to end with the bride suing the Cuban government for non-consensual sex. Yet that was the conclusion of Ana Margarita Martinez's marriage to Juan Pablo Roque, after he left their Florida residence, never to return (Bragg 1999). The mystery of his disappearance was resolved a few days later when he appeared on television broadcast from Havana, unveiling himself as an undercover spy sent to infiltrate the dissident community in the United States. This was surprising, and unwelcome, news to Martinez, who thought she had married a fellow dissident. Consequently, she filed a lawsuit about the deception against his employer. Since Cuba was not accustomed to defending itself in Florida's courts, it did not contest the suit, and Martinez was awarded millions of dollars in damages. Part of her case was based on her claim that the fraud meant that she did not give valid consent to sex with Roque. As her lawyer put it, "She would not have given her consent, had she known."

Martinez's case is controversial—is sex non-consensual when it results from someone disguising their career, allegiances and motives? This question is difficult and weighty: typically, it is appropriate to censure non-consensual sex heavily in law. Consequently, sexual consent has been a focus of research into the relationship between deception and consent, and this chapter will reflect this emphasis. But the issue also arises for other types of consent. If someone impersonates a police officer, then she lacks valid consent to entering someone's home. If someone misrepresents the risks of a clinical study, then she lacks her research subjects' valid consent.

For the purposes of this chapter, let us define non-consensual behavior as behavior for which the agent needs, but lacks, the other person's valid consent. Then we can ask two key questions. First, when does deception lead to non-consensual behavior? Second, why does it do so? These questions are orthogonal to other moral evaluations of deception. Perhaps, deception is *pro tanto* morally wrong, because it involves treating people as "mere means" (O'Neill 1985), and undermines valuable communication (Shiffrin 2014). Or perhaps deception can be benign or even salutary, protecting our privacy (Nagel 2006), smoothing social interactions (Calhoun 2000), or allowing mystery to build attraction (Buss 2005). These moral evaluations are separate from the relation between deception and consent.

Deception that is morally objectionable in itself might not lead to non-consensual behavior, and vice versa.

This chapter focuses on “permissive consent” (see this Handbook’s Chapters, *THE SCOPE OF CONSENT* and *THE NORMATIVE FORCE OF CONSENT*). Permissive consent waives a right against an interaction, releasing someone from her duty not to engage in it. So wherever rights protect realms of our lives, permissive consent allows us to share these realms. In addition to permissive consent, the English word “consent” can refer to agreements that create obligations, such as promises. Deception can also be a problem for these agreements, but while this topic is important, it will be passed over here.

For our purposes, let us define deception as intentionally causing someone to form a false belief. This definition includes paradigmatic lies — understood as deliberate assertions of falsehoods — as well as misleading statements and omissions. It excludes omissions that merely allow pre-existing false beliefs to persist. It also excludes concealment that does not create false beliefs. Currently, I am concealing from you my address, without deceiving you.

This chapter will discuss accounts of how and why deception leads to non-consensual behavior, introducing a distinction between “duty-first” and “awareness-first” approaches. Duty-first approaches begin by specifying key duties not to deceive a consent-giver. These approaches then hold that the consent-receiver lacks the consent-giver’s valid consent because she has not carried out these duties. By contrast, awareness-first approaches start with an account of key features of an interaction of which a validly consenting person must be aware. If they like, an awareness-first theorist can then appeal to this awareness to derive duties not to deceive. To be clear from the start, these approaches need not be exclusive, since it may be that there is more than one way that deception can lead to non-consensual behavior. Moreover, there is little reason to think they are jointly exhaustive, as it is plausible to think that there are other theoretical options waiting to be explored.¹ Still, it is hopefully a helpful distinction for orientating us with the topic of deception and consent.

Duty-first approaches

A duty-first approach begins with the thought that valid consent requires certain background conditions to be in place. For example, we might think of giving consent as a form of transaction. To ensure this is a “fair transaction,” the consent-receiver may need to discharge key duties towards the consent-giver (Miller & Wertheimer 2009). These duties would plausibly

include duties not to coerce the consent-giver, as well as duties to disclose information and to avoid certain types of deception. A duty-first approach then appeals to these duties in order to develop an account of when the deceiver lacks the consent-giver's valid consent to an action: when the deceiver fails to discharge one of these duties, the victim does not validly consent to the action.

This order of explanation seems to fit many people's intuitions about when a deceiver lacks valid consent since it places importance on the aetiology of false beliefs relevant to decisions to consent (Beyleveld & Brownsmith 2007: 147). If false beliefs are not caused by the consent-receiver, then the context need not be unfair. But if these beliefs result from the consent-receiver not fulfilling her duties, then the consent-receiver could lack valid consent. Placing importance on the aetiology of false beliefs is the main advantage of a duty-first approach. Its main disadvantage is that it owes us an account of the boundary between key duties not to deceive (which must be discharged for valid consent to the encounter in question) from minor duties not to deceive (which need not be discharged for valid consent to that encounter).

Determining these key duties is a complex task. The duties plausibly vary with different types of consent. For property consent, some deception seems fair game. For example, in commercial exchanges, both parties may disguise their interests and resources, in the course of arriving at a valid agreement. For example, Smith might play down her interest in leasing Jones's car, while Jones might mislead Smith about the lowest price that she would accept. Moreover, people are allowed some "puffery" when presenting the value of their property (Shiffrin 2014). This is not to say that anything goes in property exchanges. Vendors and lessors may not deceive consumers about critical features of their products. Indeed, this would frequently be fraud. Similarly, a consumer may not pay with counterfeit money. It is hard to distinguish acceptable and unacceptable deception in property exchanges, and there may be several ways of doing so. Here we encounter broader philosophical questions concerning the role for conventions in the ethics of property. Plausibly, the answers to these questions will determine much about our key duties not to deceive with respect to property consent.

What about consent to therapeutic medical treatment or participation in research? Here legal and institutional norms bear on how deception relates to consent. Professionals are obligated by institutional codes, and legally obligated by the laws of their jurisdiction. Assuming these codes and laws are adequately just, they bear on moral obligations. These

codes do not just cover deception, but also duties of disclosure. These duties are linked. If a physician must disclose a risk factor, then she is required not to deceive the patient about the risk factor, and not to conceal the risk factor from the patient. Since duties of disclosure generate both duties not to deceive and not to conceal, we should take disclosure as fundamental here. Consequently, we should start theorizing how deception relates to consent to research and therapy by thinking of duties of disclosure. (Manson & O'Neill 2003; see also the chapters in this handbook on INFORMED CONSENT; CONSENT AND MEDICAL TREATMENT; and CONSENT IN CLINICAL RESEARCH.)

For sexual consent, the issue is complicated by the fact that deception features widely in people's sexual lives (Buss 2005). The relevant philosophical debate has been framed by a distinction in Anglo-American law between "fraud in the *factum*" and "fraud in the inducement." (Feinberg 1986, following Perkins 1957: 856; Westen 2004). Fraud in the *factum* misrepresents the nature of the sexual encounter itself. The law typically conceives of this narrowly, including impersonation or disguising a sexual encounter as a medical procedure (Falk 1998). This deception would typically lead to legally non-consensual sex. All other deception is considered fraud in the inducement. A paradigm would be deception about one's romantic intentions. The thought is that if someone engages in a sexual encounter because her partner pretends to be interested in a long-term relationship, then she is not misled about the sexual encounter itself and so still validly consents to the sexual encounter. The phrase "fraud in the inducement" may suggest it covers only fraud about people's incentives. But for the *factum* / inducement distinction to be exhaustive, the inducement category must also cover deception about any "peripheral" feature of the interaction that is not fraud about one of the "core" features that constitutes fraud in the *factum*.

The *factum* / inducement distinction has received much criticism (Wertheimer 2003; Manson ms). For a start, it is hard to offer a principled basis for distinguishing facts that concern the sexual encounter and facts that do not, without this distinction yielding an over-inclusive conception of fraud in the *factum*. What should we "pack into" the description of the sexual act itself (Feinberg 1986; McGregor 2005)? Should the act be described as sex with a married man? If so, then a man pretending to be single would be committing deception about the sexual act itself. In addition, it appears *ad hoc* to apply the *factum* / inducement distinction to sexual consent, given a similar distinction would bear no weight in other domains (Schulhofer 1998; Rubinfeld 2013). For example, if one person

fraudulently offers inducements for engaging in a property exchange, then this fraud will typically leave the exchange legally void. These criticisms would target the *factum* / inducement distinction both as a principle for rape law, and as the basis for moral principles.

If we reject the *factum* / inducement distinction, what can we put in its place to distinguish e.g. key duties not to impersonate people (which must be fulfilled for valid consent to the encounter in question) from minor duties not to lie about one's income (which need not)?

One option appeals to people's expectations about the likelihood of different harms. Perhaps, an impersonator lacks their victim's valid consent because impersonation is so unexpected. Meanwhile, people lie so often about their income that detecting such a lie would not lead to a deep sense of harm. This option could yield either normative principles for framing the criminal law (McGregor 2005: 181-189), or moral principles. Indeed, it seems independently plausible to emphasize expectations if we think that morally valid consent consists in communication. What a speech-act communicates depends on conversational participants' shared assumptions (Stalnaker 2014). These assumptions will typically depend on the social and legal context.

Another option distinguishes types of information about which someone is required or permitted to share with sexual partners. For example, someone might be required to disclose STIs, but not whether they were a victim of a sex crime. In addition, it could be that we are required to disclose some information on request, but we need not provide of our own accord. This gives us an account of our duties of disclosure. We could use this account to say that a deceiver lacks their victim's valid consent to sex if and only if the deception involves failing to discharge a duty of disclosure (Lazenby & Gabriel forthcoming). Thus, we would derive norms governing deception from norms of disclosure, in the way that we do for medical and research consent.

A more radical option holds that we have duties not to deceive others about *any* feature that is materially relevant to their decisions to consent (Dougherty 2013).ⁱⁱ We might worry that alternative views implicitly rely on a notion of sex that is moralized insofar as certain features of a sexual encounter are deemed more objectively important than other features, independently of what the participants consider important (Dougherty 2013). Moreover, we may think that a "fair transaction" account suggests a radical position on deception and consent (Miller & Wertheimer 2009). Deception that is materially relevant to someone's decision-making is a paradigm of manipulating someone or treating her

“merely as a means” (O’Neill 1985; Pallikkathayil 2010). Consequently, material deception would typically bring about an unfair context for consent. This option is radical because it implies that deception about one’s income will lead to non-consensual sex when income is material to someone’s decision to consent to sex. Consequently, we would need to either revise down our estimate of the moral and legal importance of consent (Rubenfeld 2013), or revise up our estimate of how morally wrong deceiving someone into sex is (Dougherty 2013) and how stringently the law should penalize it (Larson 1993; Herring 2005).

Awareness-first approaches

Awareness-first approaches start with an account of the awareness that someone must have of an interaction, in order to validly consent to it. If they like, these approaches can then use this account to derive our duties of disclosure and our duties not to deceive.

Consequently, a principal benefit of an awareness-first approach is potentially that it provides us with a framework to determine the nature of these duties. To complete this task, we would need to consider other interests, such as consent-receivers’ interests in privacy or in avoiding various costs of disclosure. But once we have factored in these interests, we have an underlying rationale for the existence and shape of these duties: these duties’ goal is to provide consent-givers with the awareness necessary for valid consent.

The principal cost to an awareness-first approach is that it does not place any intrinsic importance on the aetiology of false beliefs or ignorance when determining whether these lead to non-consensual behavior. If a consent-giver’s false belief by itself is enough for an absence of valid consent, then it would not matter, with respect to the validity of the consent, whether this false belief has arisen by an innocent mistake, by non-disclosure or by deception. This will strike many as counterintuitive. An awareness-first theorist could sweeten this pill by allowing that other moral phenomena are sensitive to this aetiology. For example, the consent-receiver’s *blameworthiness* for non-consensual behavior could depend on whether she knew of the false belief, and in turn this knowledge may depend on the belief’s aetiology. Similarly, this knowledge may influence how morally problematic the interaction is in other respects; for example, the knowledge may influence the extent to which someone is being manipulated, and manipulation may be wrong independently of considerations of consent.

How do we determine what awareness is necessary for valid consent? To answer this, we might start with the thought that the validity of consent depends on the voluntariness of the consent. In paradigmatic voluntary agency, an action is imputable to an agent because the action is caused and rationalized by the agent's conative and cognitive states. Coercion weakens the causal links between her conative states and the action—the agent does not otherwise want to perform the action except to avoid the coercer's threats. Consequently, to the extent that someone's action is coerced, the action did not flow from her own voluntary agency. According to one view, for consent to be morally valid, it must meet a threshold for voluntariness; so, coercion could invalidate consent by making it insufficiently voluntary.ⁱⁱⁱ

We might think that similarly, deception can invalidate consent by making it insufficiently voluntary (McGregor 2005: 181). Just as coercion weakens the connection between the agent's conative states and her action, false beliefs and ignorance could weaken the connections between her cognitive states and the action. If the agent is unaware of features of her action, then she does not intend these features. So if we were to describe what the agent does intentionally, we would not mention these features. In the case of consent, we might say that it is only to the extent that the agent is aware of an action's features, she voluntarily consents to *this action*. In this way, we might hold that to the extent that someone is misled about an action, the less willingly she engages with this act, and she is wronged to the extent that her will is not implicated in the act (Archard 1998: 50).

Along these lines, we might classify features of an act, according to how undesirable the consent-giver finds these features. If a feature of an action only puts her off the action in a limited range of circumstances, that feature would be a reasonably weak deal-breaker. But if the feature put her off the action in nearly all circumstances, then it would be a fairly strong deal-breaker. Consequently, we could determine the voluntariness of a decision to consent to an action, according to the strength of any deal-breakers that are being concealed (Manson 2017).

Alternatively, for the purposes of determining the validity of consent, we could distinguish gradations of voluntariness in terms of the means by which the deception is achieved, e.g. whether it involves promises, threats, harms or non-benefits (Feinberg 1986). This approach leaves us with the further question of how exactly a lack of awareness makes one's consent involuntary? This seems as substantive as a prior question

of how a lack of awareness could lead to an absence of valid consent. So it is unclear how much progress has been made with this approach.

Instead of appealing to voluntariness, we could focus on the *content* of someone's consent. For example, we might claim that for some types of consent, someone validly consents to an encounter only if they intend to consent to that encounter. This claim follows straightforwardly if consent consists in a mental attitude (Hurd 1996; Alexander 1996, 2014. See the chapter in this handbook, WHAT IS CONSENT?). But it could also follow if we think that such an intention is necessary (although not sufficient) for some types of consent to be valid. Sexual consent might be a plausible example. If an intention to consent to an encounter is necessary for validly consenting to the encounter, then we should consider the content of such an intention. This will determine whether the actual encounter that took place is one such that the consent-giver intended to consent to that encounter (rather than an entirely different, merely hypothetical, encounter) (Dougherty 2013; Manson ms).

Let us start with the fact that an intention is a propositional attitude. Suppose Jane intends that John cuts her hair. The content of Jane's intention is the proposition that John cuts Jane's hair. In attributing this intention to Jane, we need not suppose that this is the only relevant intention of Jane's. For example, we should also attribute to Jane an intention whose content is the proposition that the man in front of her cuts her hair. In this way, our attributions of intentions are similar to our attributions of other mental states, such as belief. When the haircut begins, Jane will not only believe that John is cutting her hair; she will also believe that the man in front of her is doing so.

At this juncture, we must say more about the nature of these propositions that constitute the content of these mental states. Consider a case of impersonation. Suppose Brian impersonates Robert Redford. As Brian is standing before Alice, she intends to have sex with him, thinking he is Robert Redford. It becomes tricky to characterize Alice's attitudes because there is a man in front of her, and she takes herself to be consenting to that man; but this man is Brian and Alice does not take herself to be consenting to sex with Brian. We could say that the contents of Alice's intentions are "mental descriptions," (Hurd 1996). For example, Alice is consenting to sex with someone under the mental description of "Robert Redford." Since this description picks out a man distinct from Brian, we can point to a key intention of Alice's that does not include Brian in its content. Consequently, we can point to this mismatch between the event

in the world, and the content of Alice's intentions, in order to explain why she did not validly consent to the event (Hurd 1996).

Although we theorists would be forced to use language to characterize these mental descriptions, we should not assume that the mental descriptions themselves are linguistic entities. In other words, we need not think that mental descriptions are either natural language sentences that are consciously brought before the consent-giver's mind, or sentences in a language of thought. Indeed, there is no reason to think that a natural language like English is sufficiently expressively rich that we can always adequately articulate the relevant mental description using this language (McDowell 1984).

The idea that Brian does not fall within the content of Alice's intention has sometimes been put in terms of a distinction in the philosophy of language between *de dicto* and *de re* attributions of attitudes. Heidi Hurd takes this approach in a passage that has been so influential that it is worth quoting in full:

Suppose that Alice slept with a man she thought was Robert Redford. He was not. Did she consent to sleep with him? The answer turns on whether her intention was *de re* or *de dicto*. If she intended to sleep with "that man"—whatever his name, profession or personal characteristics—then her intention was *de re*. Insofar as she slept with "that man," she did what she intended to do, and hence, she consented to his actions.

If, on the other hand, she intended to sleep with "the Sundance Kid of movie fame," then her intention was *de dicto*. Insofar as the man with whom she slept was not the famous Robert Redford, she did not intend, and hence, did not consent, to engage in sexual intercourse with him.

Most intentions are *de dicto*, not *de re*. That is, most intentions are held under discrete descriptions. (Hurd 1996: 127; See also Liberto forthcoming)

Hurd is right to focus our attention on the exact content of Alice's consent. Using square brackets to indicate the "mental description" under which Alice conceptualises her intention, there is an important difference between these two intentions:

(1) Alice intends that [she has sex with that man, whoever he is]

(2) Alice intends that [she has sex with the Sundance Kid of movie fame]

However, we should not follow Hurd in characterizing (1) as a *de re* intention, and (2) as a *de dicto* intention. Instead, if we use the *de re / de dicto* terminology in the standard way that it is used in the philosophy of mind and language, then both (1) and (2) are *de dicto* intention ascriptions. This is because both characterize the “mental description” of Alice’s intention in terms of concepts that Alice herself must possess (Burge 1977). By contrast, *de re* ascriptions of attitudes merely indicate relations of acquaintance. Consequently, these ascriptions can be made when someone is acquainted with the subject matter of her attitude (Lewis 1979). We can make such an ascription in the case of Alice, since she is acquainted with someone—Brian:

(3) Of Brian, Alice intends that she has sex with him

Unlike (1) and (2), there are no square brackets in (3), since the *de re* intention ascription does not characterize the “mental description” under which Alice intends sex. Consequently, *de re* ascriptions need not employ terms that express concepts employed by the subject herself. For example, (3) does not imply that Alice is thinking of Brian under the mental description “Brian.” Indeed, (3) can be true even if Brian has falsely said his name is “Robert.” The point that is crucial for our purposes is that simply because Alice is directly acquainted with Brian as her sexual partner, (3) will be true. Consequently, (3) will be true regardless of whether (1) is true or (2) is true. But if (2) and (3) are both true, then Alice does not validly consent to sex. As a result, (3) is insufficient for Alice to validly consent to sex with Robert. In other words, a *de re* intention is insufficient for valid consent. This point generalizes widely: because (conscious) sex involves acquaintance, *de re* sexual intentions are ubiquitous, and so will always be present when deception leads to non-consensual sex.

Again, it is important to bear in mind that a consent-giver will typically have multiple intentions. For example, Alice will no doubt also have this intention:

(4) Alice intends that [she has sex with the man standing in front of her]

Indeed, we need to attribute to Alice a *de dicto* intention that employs an indexical concept like “in front of her,” in order to explain e.g. her physical interactions with this man (Lewis 1979). Moreover, it is a *de dicto* attitude ascription like (4) that grounds the *de re* attitude ascription of (3). In other words, it is true that of Brian, Alice intends to have sex with him because Alice intends that [she has sex with the man standing in front of her]. As we have already noted, this attribution of the intention in (4) does not exclude attributions of more specific intentions like (1) and (2). Clearly, Alice will have intention (4) if she has intention (1). But she will also have intention (4), if she has intention (2). (Compare the analogues of predicative belief attributions to (1), (2) and (4): clearly, Alice has the unspecific belief that she will have sex with the man standing in front of her.) So if we assume that Alice does not give valid consent if Brian has impersonated Robert Redford and (2) is true, then we should conclude that (4) is insufficient for Alice’s giving valid consent. To put that point in the language of “mental descriptions,” it is not enough for Alice to validly consent to sex with Brian that Alice consents under *some* description to sex with Brian.

Assuming Alice does not validly consent to Brian, how should we explain this? It is natural to look for mental descriptions, such that if someone gives sexual consent under one of these descriptions, and the description is not true of her sexual partner, then she does not validly consent to the sex. This is a difficult and substantive issue. One promising approach is to consider the question of which rights we have in the first place. One answer is that we only have coarse-grained rights not to have sex with someone (which we can waive for any reason at all), but lack fine-grained rights not to have sex with someone-with-a-headache (Liberto forthcoming). On this view, if Jones intends to consent to sex with Smith under the description “Smith,” then this intention could constitute valid consent, even if Smith has a headache and Jones does not intend to consent to sex with Smith under the description “Smith with a headache.”

Alternatively, we could adopt a radical view that focuses on the features of the sexual encounter that are material to someone’s decision to consent to sex. Let us define a “deal-breaker” as any feature of an encounter, such that the consent-giver is opposed to consenting to the encounter, given that it has this feature.^{iv} Now a consent-giver does not intend to consent to any encounter that has a deal-breaker. So if such an intention is necessary for valid consent to this encounter, then the following result holds: if an encounter only occurs because the consent-giver was deceived about a deal-breaker, then she does not validly consent to that encounter (Dougherty 2013). This result will strike many people as coun-

terintuitive, since it entails that even deception about someone's income could lead to non-consensual sex, if a certain income was a deal-breaker for the consent-giver. The result may be more palatable if we allow that some forms of non-consensual sex are only minor wrongs, e.g. when the sex imposes only minor harms on the consent-giver (Boonin ms).

In our discussion, we have been looking at the content of a consent-giver's intentions. But in closing, we should note that there is another option: we could look at the content of their expressions of consent, and consider when and how deception might result in this content failing to include the interaction that actually takes place. This option will seem particularly appealing if we think that communication is necessary for valid consent (Dougherty 2015). But it is an option that has received little attention to date, and so for now must remain an avenue for future research.

Conclusion

This chapter has discussed different accounts of the relation between deception and consent, organized around a novel distinction between duty-first and awareness-first approaches. While analytically distinct, these approaches need not be competitors, as there may be more than one way that deception relates to consent. The approaches have different advantages and disadvantages. Duty-first approaches start with an independent account of key duties not to deceive, and then use this account to say that when these duties are not met, the deceiver lacks her victim's valid consent to the action in question. These approaches have the advantage of making the aetiology of false beliefs intrinsically morally relevant—they claim that there is no valid consent only when the false beliefs are caused by the consent-receiver failing in her duties. But a disadvantage of duty-first approaches is that they now owe us an account of what these key duties are. This disadvantage is not faced by awareness-first approaches. These approaches begin by specifying an independent account of the awareness necessary for valid consent. This opens up the possibility of appealing to this account in outlining the rationale for our duties not to deceive. But the approaches come with the cost of failing to make the aetiology of false beliefs intrinsically relevant to whether consent is valid.

Duty-first and awareness-first approaches both need to choose which features of an encounter are such that deception about these features will lead to an absence of valid consent to this encounter. They are choosing between similar options. For example, they could focus narrowly on some special features, or broadly include any features that the consent-giver considered materially relevant to her consent. But while both duty-

first and awareness-first approaches share these options in common, their motivations for choosing between these options may differ. On duty-first approaches, we would fundamentally be concerned with identifying the duties that one must fulfil in order to create an e.g. adequately fair context for a consent transaction. This would frame our choice as one to be guided by concerns of fairness. On an awareness-first approach, we would fundamentally be concerned with identifying true beliefs that one must have either for consent to be voluntary, or for the consent's content to include the actual encounter. In this way, the distinction between duty-first and awareness-first approaches is not only useful for characterizing the internal justificatory structure of a view of consent. It also helps make clear how different approaches should go about determining which types of deception lead to non-consensual behavior.

References

- Alexander, L. (1996) "The Moral Magic of Consent II," *Legal Theory* 2: 165-174
—2014. "The Ontology of Consent," *Analytic Philosophy* 55: 102-113
- Archard, D. (1998) *Sexual Consent*, Oxford: Westview
- Beyleveld D. & Brownsword R. (2007) *Consent in the Law*, Oxford: Hart Publishing.
- Boonin, D. Manuscript "Sex, Lies and Harm: Dougherty on Deceptive Seduction"
- Bragg, R. (1999) "Ex-Wife is Suing Cuba over a Spy's Deception," *New York Times*, August 15 accessed on February 1, 2013 at <http://www.nytimes.com/1999/08/15/us/ex-wife-is-suing-cuba-over-a-spy-s-deception.html>;
- Burge, T. (1977) "Belief *De Re*," *The Journal of Philosophy* 74: 338-362
- Buss, S. (2005) "Valuing Autonomy and Respecting People: Manipulation, Seduction and the Basis of Moral Complaints," *Ethics* 115: 195-235
- Calhoun C. (2000) "The Virtue of Civility," *Philosophy & Public Affairs* 29: 251-275
- Chadha, K. (2016) *Sexual Consent: Sketching a Liberal Account*. Thesis submitted in partial fulfilment of the requirements for the degree of BPhil in Philosophy, University of Oxford
- Dougherty, T. (2013) "Sex, Lies and Consent," *Ethics* 123: 717-744
- Falk, P. (1998) "Rape by Fraud and Rape by Coercion," *Brooklyn Law Review* 64: 39-149
- Feinberg, J. (1986) "Victims' Excuses: The Case of Fraudulently Procured Consent," *Ethics* 96: 330-345

- Herring, J. (2005) "Mistaken Sex," *Criminal Law Review* : 511-524.
- Hurd, H. (1996) "The Moral Magic of Consent: I," *Legal Theory* 2: 121-146
- Jubb, R. (Forthcoming) "Consent and Deception," *Journal of Ethics and Social Philosophy*
- Larson, J. (1993) "Women Understand So Little, They Call My Good Nature 'Deceit': A Feminist Rethinking of Seduction," *Columbia Law Review* 93: 374-472
- Lazenby, H. & Gabriel, I. (Forthcoming) "Permissible secrets," *The Philosophical Quarterly*
- Liberto, H. (Forthcoming) "Intention and Sexual Consent," in M. Smith & U. Heuer (eds.) *Philosophical Explorations*, Special Issue
- (Ms) "Moral Debilitation"
- Lewis, D. (1979) "Attitudes *De Dicto* and *De Se*," *The Philosophical Review* 88: 513-543
- Manson, N. (2017) "How Not to Think about the Ethics of Deceiving into Sex," *Ethics* 127: 415-429
- (Ms) "Why Deception Does Not Undermine Consent"
- Manson, N. & O'Neill, O. (2007) *Rethinking Informed Consent*, Cambridge: Cambridge University Press
- McDowell, J. (1984) "*De Re* Senses," *Philosophical Quarterly* 34: 283-294
- McGregor, J. (2005) *Is it Rape? On Acquaintance Rape and Taking Women's Consent Seriously*, Aldershot: Ashgate
- Miller, F. & Wertheimer, A. (2009) "Preface to a Theory of Consent Transactions: Beyond Valid Consent," in F. Miller & A. Wertheimer (eds.), *The Ethics of Consent*, Oxford: Oxford University Press
- Nagel, T. (2006) "Concealment and Exposure," *Philosophy & Public Affairs* 27: 3-30
- O'Neill, O. (1985) "Between Consenting Adults," *Philosophy & Public Affairs* 14: 252-277
- Pallikkathayil, J. (2010) "Deriving Morality from Politics: Rethinking the Formula of Humanity," *Ethics* 121: 116-147
- Perkins, R. (1957) *Perkins on Criminal Law 1st ed*, New York: Foundation
- Rubinfeld, J. (2013) "The Riddle of Rape-by-Deception and the Myth of Sexual Autonomy," *Yale Law Review* 122: 1372-1443
- Schulhofer, S. (1998) *Unwanted Sex*, Cambridge MA: Harvard University Press
- Shiffrin, S. (2014) *Speech Matters: On Lying, Morality and Law*, Princeton NJ: Princeton University Press
- Stalnaker, R. (2014) *Context*, Oxford: Oxford University Press
- Tadros, V. (2016) *Wrongs and Crimes*, Oxford: Oxford University Press

Wertheimer, A. (2003) *Consent to Sexual Relations*, Cambridge: Cambridge University Press

Westen P. (2004) *The Logic of Consent*, Aldershot: Ashgate

ⁱ One option, which may turn out to be an alternative to both approaches, is the idea that deception leads to an absence of valid consent for an encounter when the deception has a sufficiently adverse effect on the meaning and value of the encounter (Tadros 2016).

ⁱⁱ For criticism of this option, see (Manson 2017; Jubb forthcoming; Lazenby & Gabriel forthcoming; Liberto forthcoming).

ⁱⁱⁱ For (to my mind, decisive) criticism of this view of why coercion debilitates consent, see (Liberto ms). See also (Tadros 2016).

^{iv} As a flawed rule of thumb, we might employ counterfactuals: for most deal-breakers it will be the case that, were the consent-giver aware of the feature, then she would not have decided to consent. But this counterfactual test is only good as a rough and ready heuristic and faces counterexamples—see (Chadha 2016; Tadros 2016). To canonically state what a deal-breaker is, we should not consider counterfactual mental states, but instead focus on the actual content of the consent-giver's actual mental states (Dougherty 2013).