

Moral Risk and Communicating Consent

Renée Jorgensen Bolinger

In addition to protecting agents' autonomy, consent plays a crucial social role: it enables agents to secure partners in valuable interactions that would be prohibitively morally risk otherwise. To do this, consent must be observable: agents must be able to track the facts about whether they have received a consent-based permission. I argue that this morally justifies a consent-practice on which communicating that one consents is sufficient for consent, but also generates robust constraints on what sorts of behaviors can be taken as consent-communicating.

Keywords: consent, moral rights, communication, ontology of consent

I. Setting the Stage

An agent's rights restrict how it is permissible for others to act. The fact that a mug is yours prohibits me from taking it, but you can give me permission to borrow it, and if you do, you make it the case that I would not wrong you (nor trespass your rights) were I to borrow the mug. There are many considerations other than your consent which could make it all-things-considered permissible for me to borrow your mug, for instance if doing so were necessary to save someone's life, and would only mildly inconvenience you. But when the only consideration bearing on whether I may borrow the mug is whether you allow me to do so, I require a *consent-based permission*. If I lack such a permission, taking the mug would trespass your property rights, wronging you. So if I need a mug and you want to let me borrow yours, it is important that you be able to let me know that you have given me permission.

In everyday discussions, the term 'consent' can be used in a wide array of ways. It isn't clear that all these uses track a single phenomenon, and even if they do, we might reasonably worry that our social practices regarding consent are defective, not perfectly corresponding to the moral profile of consent. So this paper is not concerned with everything we call 'consent'; it is instead wholly focused on understanding what is necessary for issuing genuine moral consent-based permissions. Much of the philosophical and legal attention given to this question focuses on a few high-stakes contexts: consent to sex, to medical procedures, or to legally binding contracts. These are some consent-based permissions, but there are also more mundane cases, such as agreeing to swap seats on a flight, inviting someone onto your property, or purchasing a coffee, and a good account of consent should be serviceable in these domains as well.¹

¹ Or, alternatively, it should justify treating what appear to be different uses of the same normative power, consent, as in fact distinct phenomena.

Actual attempts to consent may be defective in various ways; the agent might lack crucial information, or be coerced, or lack the authority to give permission. I will abstract away from these various failings to fix on *morally valid consent*, which stipulatively refers to consent that succeeds in yielding a consent-based permission. I will argue that for consent to fulfil its function in managing agents' moral risk in the wide array of cases for which it is needed, communicative behavior must ordinarily be sufficient for consent; the intention to consent is not strictly necessary. I'll devote the bulk of the paper to developing an argument from the function of consent to this conclusion; the remainder I'll dedicate to articulating an adequately constrained account of what constitutes communicating consent, and anticipating some objections.

Consent transforms an action that would otherwise wrong an agent A by trespassing her rights into a permissible act, consistent with her rights. How? Surely through A's exercising her moral agency in a very particular way, shifting the boundaries of her rights, so to speak. This implies that there must be a tight connection between consent and moral agency: whether A consents must be agency-involving. Call this thought AGENCY, since we'll want to refer back to it later. It is commonplace to assume that AGENCY implies that consenting is necessarily done *intentionally*, requiring a particular internal condition of the agent: an attitude or decision to the effect that she has no complaint against X.² I'll call such a condition 'internal assent'.

Consent also changes others' reasons for action. A's consent changes what it is permissible for B to do, and thus also changes the obligations of third parties. If B takes A's mug without her consent, other members of the community have some reason to help A recover her property. But if A consented, we lack this reason to interfere. To perform this function, it seems that consent must be OBSERVABLE: the recipient (and relevant third parties) must be able to track whether consent has been given. This pushes us to take communicative behavior to be sufficient for consent, and so is in some tension with AGENCY as commonly understood.

Advocates of attitudinal accounts (also called 'subjective' or 'mental state' views) take the consent-giver's internal assent to be necessary and sufficient for consent, and so are prepared to deny that consent must be observable. They hold that though one typically uses a behavior or speech act to communicate that one has consented, the communication itself has no effect on consent-based permissions and obligations: it is just defeasible evidence of consent.³ To explain why whether A has

² The commitment to intention is one of few nearly universal commitments in the consent literature; the assumption that consent must be intentional is so widespread, in fact, that it appears as a premise in arguments for taking intent to be a necessary condition for valid consent. For example, Dougherty (2015, 229) writes "Since consent must be intentional, everyone should agree that an intention is necessary for morally valid consent." Liberto (2017) argues, contra Dougherty (2013), that to consent one need not intend *every significant aspect* of the thing to which one consents, but she is happy to grant that consent must be intentional.

³ Several theorists have advocated attitudinal accounts, offering a variety of different candidates for the nature of the requisite attitude. For instance, Hurd (1996) characterises consent as 'intending that B do X', Alexander (1996, 2014) as A 'deciding that she has no moral complaint against B doing X', Ferzan (2016) as 'willed acquiescence', Hickman and Muehlenhard (1999) and Westen (2004) as 'acquiescing to B's Xing.' Husak (2006) also advocates a attitudinal view, but is noncommittal about the relevant attitude.

communicated that he consents makes such an intuitive difference to the permissibility of B's action, attitudinal theorists (see especially Hurd 1996; Alexander 2014) distinguish sharply between B's moral permissions and B's blameworthiness. On these views, while my uncommunicated consent can make it permissible for you to borrow my mug, you would still be blameworthy if you borrow it without at least having good reason to believe that I consented to your doing so. (Similarly, if I communicate consent without in fact internally assenting, you are not morally permitted, but you would not be blameworthy if you did borrow it.)

Theorists who focus more on the transactional nature of *giving* consent take OBSERVABILITY to be quite important, and emphasise that the principal value of consent is in its role in facilitating certain interactions. They note that consent is similar to promising in its ability to alter agents' permissions and obligations, and so plausibly if promising requires communication, consent must as well. They model consent as "an act rather than a state of mind",⁴ which requires the performance of some observable act of authorisation. There are two dominant strains communication-centered view: performative accounts, and hybrid accounts. Several advocates of *Performative* accounts analyse consent on the model of speech acts: for a given action to be a *consenting*, A must perform it with the intent that B recognise, by this performance, that A has given B a consent-based permission.⁵ Other authors are less committal about the internal structure of the communication of consent, but similarly require that the consenter intentionally engage in a behavior aimed at communicating that she consents.⁶ Advocates of *hybrid* accounts hold that A consents to B's Xing iff A has the relevant internal attitude toward B's Xing, *and* communicates that to B.⁷

There is a deep tension in communicative accounts as currently developed. Internal states like intentions can vary independently of an agent's apparent communicative behavior. So, if *intent to communicate consent* is a necessary condition for consent, either on its own or as a constitutive part of the performance of consenting, the facts determining whether B has *actually* received a consent-based permission are still not observable. So while these views are motivated by the importance of OBSERVABILITY, it is at best unclear how much they are able to capture of consent's coordinating function. Seeing this tension, several prominent authors (notably Simmons (1981), Feinberg (1986a), and Wertheimer (2003, 144)) argued that conditions internal to the consenter (A's unexpressed

⁴ Archard (1998 *Sexual Consent* p.4)

⁵ Sherwin (1996) and Cowart (2004) draw on Austin, Strawson and Searle to build a speech-act analysis of consent, requiring intent as one of several conditions for performance of the relevant act. Manson (2016) analyses consent as akin to Raz's 'authoritative commands', and includes A's reflexive intention as a necessary element. Theorists offering performative views differ with respect to how explicit the communicative behavior must be, whether it must be conventional, and whether uptake is also a necessary condition for valid consent.

⁶ For example Owens (2011; Hartogh 2011).

⁷ For instance Archard (1997), Dougherty (2015), and Gruber (2016). Tadros (2016 *Wrongs and Crimes*; cf. Dougherty (2018b)) defends a hybrid view that requiring that A *attempt* to communicate to B, but not requiring uptake. Depending how we understand the relevant internal attitude—if intending to *communicate* that one is issuing a consent-based permission counts—the division between performative and hybrid views is quite tenuous (Alexander (2014) also makes this point), and at least David Owens and Emily Sherwin also should be considered advocates of hybrid views.

attitudes and intentions) cannot be necessary for consent-based permissions. However, these theorists have generally given up claiming that their resulting accounts are of *consent*. Feinberg cast his focus as analysing facts about the *responsibility* of B (which is easily interpreted as B's blameworthiness)⁸, while Wertheimer (in Miller and Wertheimer 2009) argued that consent requires intent, but is not necessary for a 'valid consent transaction', and the latter rather than the former are the source of consent-based moral permissions.⁹ Troublingly, though, divorcing the analysis of what is necessary to *give* consent from what is necessary to *receive* consent-based-permissions in this way yields the result that morally valid consent is no longer necessary for morally valid consent-based permissions.

Both the attitudinal approach and the communicative alternatives ultimately sacrifice OBSERVABILITY in order to secure AGENCY. By making consent depend on A's internal states—specifically his intentions—these views render it opaque to recipients and third parties whether consent has genuinely been given in a context. Whether this is troublesome depends on what consent is supposed to do; if in fact OBSERVABILITY is not core to consent, there is no great problem here. I am skeptical whether we could settle this question by considering our intuitions about hypothetical cases (which are murky at best, and often track theoretical commitments). Instead, I hope to make progress by considering what it is that rights-shifting powers like consent function to do, and constraining our theorising about the nature of consent by what fits best with the functional role it must play. My goal is not to give a new theory of consent, so much as to motivate and outline a way existing accounts can and must be supplemented to secure the full functional profile of consent.

2. The Coordinating Function of Moral Rights & Powers

Perhaps the most well-recognized function of consent is to facilitate the autonomy of the consenting agent by enabling her to alter her rights when that is what she genuinely desires to do. If we focus narrowly on agents' autonomy-interests in unilaterally deciding whether a given action wrongs them, the attitudinal views of consent look very attractive. In the words of one of its proponents, "If we think that what we are protecting is autonomy, then that autonomy is best respected by recognizing that the consenter has it within his or her power to allow the boundary crossing simply by so

⁸ If Feinberg (1986a, chapter 22) is read as giving a positive account of *consent*, what he proposes is very similar to what I advocate here: a moral permission given by A's communicative act which insulates B from risk by shifting the responsibility for their action to A. But he does not articulate how communication must be constrained in order to secure the appropriate moral profile for consent.

⁹ Miller & Wertheimer offer the 'fair transaction model' of consent-based permissions, on which B is morally permitted to proceed on the basis of A's consent—but in the absence of A's consent—if B has treated A fairly and responds in a reasonable manner to what B reasonably believes is A's token or expression of consent. They recognise that this is the barest outline of an account, and needs filling in, particularly in specifying what it is to treat A fairly and reasonably believe something to be an expression of consent. The account I will offer is very much in the spirit of this approach, but provides a way to fill in some of these details more concretely.

choosing. No expression is needed.”¹⁰ But this is neither consent’s only function, nor even the only way that consent serves the giver’s autonomy. True, consent allows an agent to change her rights, altering normative facts ‘by her mere say-so’, but this serves more purposes than simply deciding whether an action wrongs her. It also enables agents to participate in close interactions, of significant moral value, that would be either impermissible or prohibitively morally risky without it. Many of our moral rights—rights to property, to bodily integrity, to privacy, etc.,—keep others at a certain moral distance. But there are valuable relationships and projects that are not possible at this distance, so the interest we have in these values gives us reason to want to be able to shift the moral boundaries created by our rights, for some people, sometimes.¹¹ And because our rights not only define and protect a sphere for autonomous agency, but also assign responsibility for costs incurred when agents’ wills collide (whether culpably or by innocent mistake), we have an interest not only in making certain boundary-shifts, but in assuring others that we have done so.

Consent facilitates this, securing for agents the freedom to permissibly engage in a wider range of actions, projects, and relationships, than would otherwise be open to them. It structures and sets the terms for agents’ moral relationships, and thus in many ways is the interpersonal, moral, analogue of public law. Having something to perform this function is an important part of moral life in a social context, enabling agents to coordinate their behavior and manage the moral risks of close interactions. In this analogy, consent facilitates the giver’s autonomy by positioning her as law-giver, articulating the rules for others’ permissible interactions with her. The properties consent must have, if it is to fulfil this function, also strongly resemble several of the features Fuller identified as ‘virtues of public law’: it must be prospective, it must be observable, it must be relatively stable (not shifting undetectably), and there should be congruence between the moral permissions *given* through consent and those *received* through consent.¹² An uncommunicated internal attitude is not equal to

¹⁰ Ferzan (2016). Conceptualising consent as autonomous authorisation is perhaps especially attractive if one attempts to determine what consent is by asking what the difference is between consensual sex and rape. Particularly among legal theorists, it is somewhat common to conceptualize the principal wrongfulness of rape as the violation of the interests that are not threatened by consensual sex, and work backwards to fill in what the interests protected by consent must be. This strategy is apt to lead us astray; it may well be that the wrong of rape is more complex, and involves violating more rights, than merely non-consensual sex. By way of example: a gift can be a consensual transfer of material goods from one party to another, and if this transfer occurs without consent, it is theft, a non-consensual taking of someone’s property. The harm involved in a theft, and the wrong thereby done, varies significantly depending on what was taken, when, whether it was done under conditions that make the victim feel generally unsafe or subject to ongoing threat, whether it impacts her future plans or present projects (did it destroy her retirement nest-egg? Did it rob her of something she’s sacrificed a great deal to secure?) etc. All of these harms are relevant to understanding the wrong of the theft, but it would be a mistake to make the fact that it was a non-consensual taking bear all the normative weight, or to start our theorising about the nature of gifts as whatever protects all the interests violated by a theft.

¹¹ Shiffrin (2008) argues that these facts underlie the moral force of both promising and consent, a point emphasized as well by Watson (2009), Owens (2012), and Dougherty (2015), among others.

¹² Fuller (1969, 106)’s full list contains eight requirements: law must be (1) general, (2) publicly accessible, (3) prospective rather than retrospective, (4) clear, (5) non-contradictory, (6) possible to satisfy, (7) stable, and (8) there must be congruence between what the law requires and what is enforced. My thanks to [redacted] for suggesting this framing.

this task. To see this additional value of consent—and why to OBSERVABILITY is key to achieving it—imagine we were to attempt to pursue close relationships without consent.

World Without Consent. Let's stipulate that actions that would be serious rights-transgressions if unwelcome (e.g. a sexual encounter or surgery) are permissible, and do not violate A's rights, if in fact A stably and autonomously desires them to occur.¹³ An agent (let's call him Arnold) might request that Beth engage in these activities with or for him, but doing so is morally risky for Beth. Suppose Arnold asks Beth to perform a surgery on him that, while not strictly necessary, would probably improve his life. Also suppose that Beth is conscientious: she'd like to help, but prioritizes avoiding trespassing the moral rights of others. We are often bad judges both of what we actually desire and of what a planned action will be like. Despite Arnold's sincere request, he may discover in the course of the activity that he was mistaken; he did *not* all-things-considered desire the surgery; in which case (in our hypothetical world-without-consent), by performing it Beth will have done a serious moral wrong. And this is still close to a best-case scenario for Beth; she has no guarantee that Arnold is even sincere. Further, because members of a moral community can be called upon to enforce each others' rights by restraining or exacting compensation from violators, if she performs the surgery Beth risks not only having done a wrong (the primary risk) but also owing compensation for it (the secondary risk).¹⁴ Beth wrongs no one if she declines, but if she performs the surgery, there is a real chance that she will turn out to have violated Arnold's rights. Any assurances Arnold can offer now are at most evidence reducing the epistemic probability she should assign to these risks eventuating. Lacking any comparably weighty moral reason in favor of running the risks, Beth might reasonably pass on such a high-stakes moral gamble, leaving him without a willing surgeon.

Consent can change this equation by allowing Arnold to “own” the risks, insulating Beth from the costs of the gamble. In giving morally valid consent to a surgery, Arnold does not guarantee that he isn't self-deceived about his desires. He simply ensures that even if he is, performing the surgery will not have transgressed his rights, so Beth will neither have wronged him nor be obliged to compensate. Consent thus makes it possible to redistribute moral risks in a way that creates space for valuable exercises of agency in close quarters.

As already noted, consent is just one of many ‘moral powers’: ways an agent can act to alter his rights, thereby changing others' moral duties and permissions concerning him. A can transfer property rights by *giving* something to B, he can undertake new duties by *promising* to do something, and he can release B from certain duties by *waiving* his rights, *consenting*, or by doing something that *forfeits* those rights. Each of these actions changes the landscape of B's moral permissions. Each thus

¹³ Note that simply *desiring that B X* is not one of the attitudes put forward by internal attitude theorists as constituting consent. I am entertaining the possibility that the presence of such a desire makes an intrusion not a rights-violation solely to motivate the moral value of consent; I will make no attempt to claim that this counterfactual is in fact true. Actual consent clearly does a lot more moral work than this, allowing A to waive rights against an action that he nevertheless desires not to occur, such as someone he dislikes attending a party he hosts.

¹⁴ Being fully excused or non-culpable for the mistake does nothing to reduce the wrong done nor the fact that compensation is owed; on most accounts, culpability affects only B's blameworthiness and liability to punishment.

carries a risk of *error*: coordination failures in which A does not intend to exercise the power, but B treats him as if he has. In each case we could take observable behavior to suffice for the exercise of the moral power, or we could insist on a division between the facts that shift the moral permissions, on the one hand, and those that provide evidence of the shift, on the other, taking the latter to be relevant only to whether an agent is blameworthy if she errs, not to permissibility. The latter strategy is precisely the one advocated by those who take the intention to consent to be necessary for consenting.

The problem is that such a thorough division between permissibility and communication subjects agents in B's position to considerable *normative opacity*: they are unable to track the moral facts relevant to whether they have received a permission. Worse, their evidence may be misleading, since they are unable to distinguish merely apparent exercises of moral powers—which do not alter permissions—from genuine ones, which do. Normative opacity creates two problems for coordinating well-intentioned agents' behavior. The first is epistemic: agents lack the information necessary to avoid accidentally transgressing others' rights while interacting; they can only avoid *culpably* doing so. The second is distributive: given that agents can't track the changes in each others' rights, they will have to find a fair, principled way of assigning responsibility for the costs incurred when well-intentioned rights-transgressions occur. Left unaddressed, the problems resulting from normative opacity can undermine the value of the moral powers: the point of such powers is (at least partially) to enable agents to interact in valuable ways by insulating them from risk of committing serious moral wrongs; merely avoiding culpability is not sufficient for this end.

To see how this happens with consent, consider Beth again. She knows that *if* Arnold consents, it's morally safe for her to perform the surgery. The value consent provides Arnold is that he can use it to remove the moral risk from Beth's decision, and so secure the interaction he values (the elective surgery). However, if a particular internal attitude is necessary for morally valid consent, then Arnold cannot assure Beth that she in fact has moral permission. All of the actions Arnold might take to communicate his consent — signing a form, saying “please perform this surgery, I understand the risks”, etc., — could in principle be done without his actually *intending* to consent or having the necessary internal attitude. So if the internal state is necessary for consent to create a genuine moral permission, then none of these communicative behaviors can assure Beth that Arnold really has given her a consent-based permission. From her perspective, because she can't track one of the necessary elements of consent, acting on Arnold's apparent consent remains a high-stakes moral gamble: if she performs the surgery and he didn't really consent, she will have done something seriously morally wrong. She might reasonably pass on taking such a gamble, leaving Arnold unable to get the surgery he wants—precisely the pickle Arnold found himself in in the *world without consent*.

Similarly, unless third parties can reliably track whether A has consented, the moral community cannot use consent to redistribute responsibility for material risks. To secure the valuable interactions that led us to posit consent as a moral power, it must somehow bringing the normative facts into line with features agents *can* track, allowing them to successfully take ownership of the relevant moral and material risks.

3. The Possibility of Consent Without Intention

To summarize: consent must be observable to perform its risk-managing functions. At the very least, B must be able to track the facts about whether A has given her a permission to X. The trouble with the views surveyed in §1 — the reason they are unable to secure the coordinating function of consent — is that they all include a particular internal state, *A's intention to change her rights*, as a necessary element of morally valid consent. Since whether A has such an intention is not observable, these accounts leave it opaque to B whether in a given instance she has received a genuine consent-based permission or a merely apparent one, and opacity reintroduces risk. If a particular internal state is necessary for moral consent, consent is unable to perform some of the core functions that justify its existence as a moral power in the first place.

An easy way to solve problems like the one Arnold and Beth face is to hold agents accountable to a conventional mapping of facts they *can* track (communicative behaviors) onto the facts of primary interest (moral permissions). If the conventional mapping is morally justified, then these behaviors become ways of exercising the moral power. Promising works like this; agents have a moral power to place themselves under new obligations by committing to do various things, and the conventional way to indicate such commitment is to say “I promise...”, or perform a similar promise-communicating behavior. These behaviors aren't merely defeasible evidence of an internal commitment; performing them is itself a way of putting oneself under a new obligation. The core cases of promising are intentional: the promisor communicates his commitment to the promise precisely because he *intends* to place himself under a moral commitment. But to keep insincerity from undermining safe reliance for recipients, the facts recipients *can* track (communicative behaviors) must be sufficient for determining whether a promise has been made; promissory *intent* can't be strictly necessary for promissory obligation. The value of the moral power of promising cannot be secured unless we address the problems of normative opacity. So that very value also justifies the introduction of a convention to settle what constitutes promising in a social context. When a society adopts such a convention, they are not merely acting *as if* it tracks the moral permissions given by promising; adopting the convention *makes it so*, because it is a necessary means to securing the justifying value of the moral power of promising. I contend that consent works in exactly the same way.

If consent is like promising, what must our account of morally valid consent look like? For a start, it will need two parts: one to characterize the core cases, in which the agent sincerely intends to exercise the power, and one to resolve normative opacity by specifying what observable facts are in any case sufficient for an exercise of the power. Debates over the nature of promising focus on the details of the core account, because the second part is uncontroversial: apparently sincere, uncanceled communication of having promised suffices to generate genuine obligation. I've here argued only that internal assent cannot be *necessary* for consent; I've made no attempt to establish that it cannot be sufficient. Accordingly, I'll remain neutral between attitudinal, hybrid, and performative views for what analysis we should give of the core cases. My focus is on the second part: developing an account of what communicative behavior suffices for consent, consistent with both

OBSERVABILITY and AGENCY. Schematically, I contend that a full account of consent has the following structure:

CONVENTIONAL CONSENT:

A consents to B's Xing if A either (i) in fact satisfies our core analysis of consenting, or (ii) communicates that they have done so, and does not cancel this communication. B receives a consent-based permission from A just if either (i) A satisfies the core analysis of consenting, or (ii) A communicates that they consent, this is not cancelled, and is not undermined by B's best available evidence.¹⁵

In most ordinary cases, both clauses will be satisfied: A will communicate consent to B because he intentionally consents in the way specified by the core account. However, in the event that A communicates that he consents but does not satisfy (i) — e.g. because he is insincere — clause (ii) goes to work, ensuring that nevertheless B enjoys a genuine consent-based permission. Agents operating under such a convention remain unable to tell whether someone who communicates consent truly intends to do so, but this opacity is rendered irrelevant; the communication suffices to alter the moral landscape. Given that A has communicated that he consents, B no longer faces a moral gamble. Embracing this extension requires giving up the thought that A only exercises a moral power intentionally; that he consents only when that is what he aims to do. This idea won't go down easily; after all, it is natural to think that if A has a right against B's Xing, then no matter how reasonably interpretable some behavior β is as communicating consent, performing β can't change A's rights without A's intending to make such a change.¹⁶

The plausibility of taking communication to be sufficient for *consenting* depends crucially on how we fill in the details about what it takes for something to communicate consent. Hand-waiving or vague answers will not do; the account is plausible only if it preserves agents' control over whether their consent is given, and it will solve the opacity problems only if receivers have generally stable knowledge about what behaviors communicate consent. The details matter here, so the next few sections are devoted to filling in what an account must look like if it is to do the necessary moral work.

¹⁵ More formally:

A's β ing constitutes consenting to B's Xing iff either (i) β is ψ , or (ii) β communicates that A ψ s.

Instances of ' ψ ' in this schema should be replaced with the preferred basic theory of consent, which will presumably include a specification of A's attitude or intention to change his rights. Note that as stated, this leaves open whether consent and consent-based permissions can arise without communication. Substituting an attitudinal account for ψ will yield the result that they can, while substituting a communicative account will dictate that they cannot. As I have framed it here, convention is relevant only at the level of taking clause (ii) to be a way of consenting. It can be made more restrictive by requiring that the communicative behavior β be one that conventionally communicates ψ .

¹⁶ See Archard (1998), *Sexual Consent*, as well as Hurd (1996) for a development of this objection.

3.1 Preserving the connection to agency

Most importantly, though the communication of consent can be unintentional, to preserve the connection to moral agency the communicative act must be something for which A is *responsible*. A can only be held responsible for communicating consent by performing some behavior β if A could reasonably be expected to be aware that β has such communicative content. Similarly, attaching moral consequences to performing β is justifiable only if it is not unduly costly for A to avoid the behavior when he seeks to avoid the moral consequence. These facts generate two constraints on which behaviors can communicate consent:

- **PUBLICITY:** It is public knowledge in the community that performance of β communicates that the performer consents to the recipient's Xing.¹⁷
- **BURDENS:** Performing β is something that every member of the community can generally avoid without undue costs.

The upshot of these two constraints is that for a behavior to communicate consent, it must be reasonable for B to assume of any arbitrary A in the relevant community that they know β is a signal of consent, and that they could have avoided it if they didn't intend to communicate their consent.¹⁸ If a behavior β satisfies both of these, then agents can reasonably demand that those who do not wish to communicate consent avoid the behavior. There will be much more to say about what it takes to satisfy these conditions, and in §4.1 I will clarify ways in which their stringency depends on context, but for now here's a case to give you a feel for the spirit of the account: it's public knowledge in our community that saying "feel free to borrow my car for the weekend; the key is on the hook by the door" communicates consent to the addressee's borrowing of your car. It's also basically costless to avoid making such a speech. If you have made such a speech, and return from your weekend conference to discover that your flatmate did indeed borrow your car, it would be inappropriate to protest that they had violated your property right, taking your car without your consent—it would be inappropriate even if, when you made the speech, you weren't really thinking about it and didn't *mean* to communicate consent.

Why? While of course your flatmate should exercise due diligence in respecting your wishes, you also have a responsibility to avoid predictably putting them in a situation in which you do not intend to consent, but which is to them indistinguishable from one in which you do. This is what it

¹⁷ Public knowledge is a less demanding standard than Lewisian common knowledge. A proposition P is *common knowledge* in a community iff everyone in the community knows that P and knows that everyone knows that P. P is *public knowledge* just if most members of a community know P, and reasonably expect others in that community to know P. This is robust with respect to isolated ignorance: a piece of public knowledge is something that every member can be expected to know, held responsible to know, and that members generally operate on the assumption that everyone knows, but that some individuals may nevertheless fail to know.

¹⁸ A subgroup of individuals in a community can of course establish their own signals; this in effect creates a smaller 'relevant community' with localized conventions. See fn 24 for further discussion.

takes to do your part in supporting the risk-managing practice that enables all well-intentioned agents to coordinate their close interactions while minimizing errors. Failing to do so is negligent or reckless, and likely to induce errors for which you can reasonably be held responsible. Thus, even unintentionally communicating consent in this way satisfies AGENCY. Like agents who insincerely or recklessly perform promise-indicating behaviors, agents who negligently or recklessly perform (and fail to cancel) consent-communicating behaviors that satisfy the constraints cannot reasonably complain against being held responsible for moral costs resulting from their behavior. Such a complaint would only succeed if the demand that he avoid β imposes more costs on him than the value afforded by the conventional consent-practice, which would only be the case if either it was difficult to anticipate that β communicated consent (failing PUBLICITY), or it was too burdensome for him to avoid β (failing BURDENS).

Importantly, however, whether A is responsible for the costs of an error is a function *only* of whether he performed an uncanceled consent-communicating behavior β that satisfied both conditions; not whether the mistake was rational on B's evidence.¹⁹ While whether a behavior is consent-communicating is not set by A's intentions, neither is it up to the receiver's interpretation. Rather, facts about which acts communicate what contents depend on signaling dynamics at the social level (which I'll explain in more detail in the next section). We should therefore distinguish two types of error: *reasonable* errors, in which though A doesn't intentionally consent, he does induce the error by performing a consent-communicating behavior β , and *rational* errors, in which B's reasons for thinking that A consented are independent of A's performing any specific consent-communicating β . A is partially responsible for a reasonable error, and that fact is relevant to whether A suffers a wrong. By contrast, in cases where B was epistemically justified in believing that A communicated consent, but in fact A neither intentionally consented nor performed a consent-communicating β , B may be blameless for the mistake, but is responsible for the error and must compensate A all the same. To illustrate, consider

Mix-up: Doctor is handed a consent form with A's signature on it by a nurse, who says he saw A sign the form. In fact there has been some sort of mix-up: A did not sign and does not assent to the surgery.

In this case, the doctor has excellent reason to believe that A has consented. But while she would be blameless in performing the surgery, doing so wrongs A by trespassing her rights, and A is owed compensation for the wrong.

3.2 Signaling: Constrained Communication without Intention

One might object that even this view requires intention, because *communicating consent* involves *intending* to communicate that one consents. I deny this, invoking instead a model of communication on which A can communicate some informational content in an agency-involving way, but without

¹⁹ It is for this reason that I have avoided framing my account as one on which *reasonable evidence of A's intentional consent* suffices for consent. Still, what I am advocating is basically a version of this, that restricts what counts as 'reasonable evidence' to observing A perform an uncanceled consent-communicating β that satisfies PUBLICITY and BURDENS.

intending to. Because this is less familiar than the speaker-intent model of communication, allow me to briefly overview the mechanics on which this account depends (drawn from Skyrms (2010)).

A behavior β signals some informational content p when, on observing β , a receiver should raise her confidence in p . Signal strength is a measure of how much the evidential probability of p changes, conditional on observing β , relative to the prior probability that p . Signs get associated with informational contents through statistical co-occurrence patterns (the frequency with which β occurs and p is true) and reinforcement pressures across the population. Here's one way that can happen: on observing β , B responds with the state-appropriate behavior for p , and things go well. B tells others, and they react to β as if p , and receivers who react as if p when they observe β rarely suffer cost-incurring errors. Quite quickly, receivers will take β as meaning that one should assume that p and act accordingly; in short, at this point in this population, β signals that p .

Signals can be more or less conventional, depending on the arbitrariness of the connection between the sign and the content. β is a maximally non-conventional signal of p if the signal strength is weaker than the natural evidential relation between β and p . β is a maximally conventional signal of p if it is a strong signal of p , but β bears no natural evidential relation to p . Most signals are somewhere in between, arising from conventional superstructures built over natural evidential connections; for instance, how smiling behavior signals feeling positively about something. Conventional signals do not ensure their contents; just as it is possible to lie, it is possible to signal that p without p being true. If this happens too often, however, it introduces 'noise' that weakens the signal.

The informational content of a behavior—what it signals—is sensitive to background information, and can be canceled by undermining evidence. For instance, in parts of the American Midwest, a siren noise at 11am on Tuesday means the tornado alert system is being tested and is working; at any other time, it means a tornado has been sighted. However, the siren on any day can fail to signal a tornado if the siren is accompanied by an announcement that "this is only a test". A signal works by indicating that the receiver should increase her credence in the content, so when accompanied by information that blocks the inference from observing the sign to its usual content, the sign does not communicate that content in that context.

Signaling dynamics make it possible for A to communicate a content to B, without having any particular communicative intention. Importantly, what is communicated by β is determined independently of either what A intends or what a particular receiver B interprets β to mean; it is set by the signalling relations operative in the relevant community. Moreover, communication is temporally extended and context-dependent. A's performance of β will only communicate consent to B in a context c if nothing in c blocks the inference from observing the behavior to A's intending to consent. Mere performance of a typically consent-communicating act won't communicate consent when accompanied by undermining information, or followed up by cancellation.

4. Constraining the View

We have allowed that one can communicate consent unintentionally, because otherwise we are left with unacceptable normative opacity. But since the justification crucially depends on this being the best way to manage agents' exposure to moral risk while enhancing their autonomy, we have reason to make it very difficult to communicate consent accidentally. So, there are some important limitations for the conditions under which β can ground a consent-based permission.²⁰

4.1 *Limitations that arise from the justification for the convention*

First, agents must defer to their best evidence, and seek confirming evidence if the costs of error are significant and asymmetric. The moral justification for this constraint is straightforward. B can only appeal to A's communicative behavior to absolve her of responsibility for an error if she was attempting to coordinate in good faith. But if, on B's evidence, it was not rational to believe that A intended his behavior to communicate consent, then a good faith effort requires B to either assume consent has not been given, or seek more evidence. This does not mean that B must be able to *rule out* or lack *any* evidence that A lacks the relevant intention; it simply must be rational, given her evidence and the costs of error (for both her and A) for her to assume that A intends his behavior to communicate his consent.²¹ In some cases, this will require her to seek a significantly more reliable signal. Suppose that β is a moderately strong signal, but has a non-negligible chance of being sent without A's intending to consent. If there is an alternative signalling behavior α that is a *very* strong signal of consent (it's *possible*, but *highly* unlikely that α would be sent without A's intending to consent), then if the costs of error are significant and fall mostly on A, α 's easy availability precludes permissible reliance on β . For example, A's being generally friendly, seeming to flirt and appearing interested in B are jointly a weak signal that A will consent to intimate involvement with B, but the fact that B could easily obtain more conclusive evidence by asking A whether he consents makes relying on the weaker signals impermissible.²²

Second, the PUBLICITY constraint is context-sensitive in stringency. The scope of the relevant community is set by the coordinating function: interpersonal permissions require only the parties to

²⁰ These limitations nicely mirror the observations in Archard (1997), Miller and Wertheimer (2009), and Dougherty (2015) about the appropriate limits of convention. "If a convention is not generally accepted, or if there is disagreement concerning its existence, and if acting in reliance upon it may result in great harm to someone, and if it is nevertheless possible directly to confirm the belief inferred from acceptance of the convention, then someone who acts in the light of a belief inferred solely from the convention may be judged to behave unreasonably and culpably." (Archard 1997, 278)

²¹ For a detailed discussion of how one might articulate the epistemic requirements for this type of rationality calculation, see Bolinger (2018) and Guererro (2019). This requirement can either be understood as internal to respecting consent, or as an independent "duty of due diligence", à la Dougherty (2018a).

²² As I have developed the view, when it is not rational on B's total evidence to assume that A intends his signalling behaviour to communicate consent, *B does not receive a consent-based permission*. The view can be made less demanding by saying instead that B receives the permission, but would be blameworthy if she did not seek stronger evidence before treating A as consenting.

the interaction, but if the legal community will be called on to do distributive or enforcement work, then the legal community is the relevant one. The set of behaviors that are publicly known to be consent-communicating is smaller the larger and more heterogeneous the relevant community is. Since for β to be an acceptable consent-communicating signal it must be public knowledge that β communicates consent, it is likely that only highly conventionalised signals will pass muster when the community is as large and diverse as a legal/political community. Interpersonal communication is more flexible, as the parties are able to create something like a more personalised small-scale signaling convention. This means that not all cases in which A communicates consent to B will be cases in which third parties (like the legal community) can recognise that A communicated consent. In such cases, the permission-giving function comes apart from its ability to change third-party's reasons for intervention: while A might have given B a moral consent-based permission sufficient to change who is *morally* responsible for an error, that fact is not adequately public for the legal community to ascertain and so it is insufficient to change who is *institutionally* responsible. If pressed the law should assume permission is absent.²³

Third, BURDENS is also context-sensitive. The justification of the consent convention is that it secures the value of the practice for the participants. So, there's a built in limit to costs of avoidance relative to the value of the practice: the guaranteed loss of the costs of avoiding β must be outweighed by the value of being able to manage moral risks by using β to indicate consent. The costs of avoidance are not limited to the difficulty for A of avoiding β at a particular moment, but rather should be understood as what A will have to sacrifice in order to avoid β whenever she does not consent over the course of her life. As a rule of thumb, if one cannot reasonably appeal to the value of reduction in moral risk achieved by clause-ii consent to justify demanding that A sacrifice β when he does not wish to consent to X, then β fails the BURDENS constraint. The practical upshot is that it is easier to accidentally consent to everyday interactions with low costs of error (borrowing a car, say) than to activities with a higher cost of error (e.g. sexual encounters or invasive surgery). This is because in the former, the value of having an easy informal way to indicate consent is greater, and the risks of accidentally indicating consent are more easily recoverable, whereas the asymmetric costs of mistake in the latter sorts of transaction force consent-communicating behaviors to be *very* easily avoidable in order to clear the burdens threshold.

²³ A reader helpfully informed me that the practices of BDSM communities provides an interesting case study of negotiating differences between local consent-communicating conventions and those operative in the broader community. Often, consenting parties desire an experience *as of* non-consent, which is achieved by first clearly communicating the limits of interactions being consented to. Because the parties in this local context have an interest in using behaviors (like saying 'no' or 'stop!') that normally signal non-consent without their normal meanings, they establish an alternative signal, the 'safe word', for the revocation of consent during the interaction. Often parties formalise this agreement in a contract. Though not legally enforceable, the contract provides evidence of the specialised communicative convention operative in the context of the interaction, which allows legal judgments of consent to be made relative to this special context. See 'Non-Binding Bondage' (2014), *Harvard Law Review* 128(2), 713-734, for more discussion.

4.2 Problem Cases

In allowing that one can sometimes give morally valid consent unintentionally, I have offered an account that seems to open the door for some abuses. So though I am in general abstracting away from defects like ignorance and coercion to focus on morally valid consent, allow me to briefly address some situations in which one might want to say consent is absent, but worry that my account cannot yield that verdict.

Let's start with coercion: in cases where B directly threatens to harm A unless he performs β , the signal is doubly undermined. First, B is in possession of signal-undermining evidence, since B has demanded that A perform the behavior, and so has reason to think it doesn't indicate A's assent. Second, B has made it the case that avoiding β is unacceptably costly for A. Under these conditions, B cannot reasonably demand that A be made to bear the costs of error, given that he performed β : he did not fail in a responsibility to avoid putting B in a position indistinguishable from one in which he intends to consent. Similarly, in contexts where B holds power over A such that failing to β when B wants A to is especially costly, β cannot communicate consent, because B's inference from A's behavior to A's intending to consent is blocked by the knowledge that in this context, avoidance is too costly. If this is true of all consent-communicating behaviors for B's Xing, it may be impossible for A to communicate consent to B in that context.

So far so good. The next worrisome case is with agents who β while ignorant that β communicates consent, even though most people in the community do know, so the signal passes PUBLICITY. Here I will say that consent *is* present: If Adele wanders into a silent auction and raises her hand in a way that, according to the conventions operative in the room, communicates placing a bid, the auctioneer has a consent-based permission to put her down as having placed a bid—unless it is clear to him that Adele does not understand the conventions and did not intend to bid. If he can discern *that*, then he has undermining evidence and is not permitted to rely on the signal. This reflects a general hazard of travel between communities with diverging communicative practices. Agents are accountable to the conventions and practices of the signaling community they happen to be in, *even if* the agent is unaware of those conventions, so long as it is reasonable for the rest of the community to think that the agent *is* aware. It is in general a foreigner's obligation to acquaint herself with the local conventions; she can insulate herself from accidentally consenting by making it clear that she is ignorant. To say otherwise allows the interloping agent to undermine the signal as a reliable indicator of the relevant moral permission, returning the community to the high-stakes gambling problem the signalling practice exists to solve. So if the practice is sufficiently valuable, the community can reasonably demand that the interloper educate herself, announce her ignorance, stay home, or bear the costs of error, rather than destabilizing their practice.²⁴

²⁴ This is not to say that it will be *easy* or *costless* for the interloper to announce her ignorance: clearly indicating that you are unfamiliar with local norms can attract attention from people who will exploit that ignorance, or subject the stranger to hostile treatment. Here there is an important asymmetry between individuals and minority groups: a new group in a region can establish a practice, and has some claim to negotiation and accommodation from their neighbouring

Note though that Adele's situation is a worst-case scenario for unintentional consent. Many, perhaps most, consent contexts are temporally extended interactions, with intricate consent-communicating behavior that it is difficult to haplessly perform. And since consent can be canceled or revoked, A can at any time that she becomes aware of a miscommunication unambiguously revoke consent for going forward. This reduces the chances of unintentional consent doing moral work in many cases, and makes it quite unlikely for something like a surgery or sexual encounter.

4.3 *Bad Conventions*

On this view, the principal function of the moral power *consent* is to coordinate agents' behavior in a risk-managing way, and this function is fundamentally social. What it takes to exercise the power can be somewhat flexible, but it must be observable: we must be able to use the facts about whether a consent-based permission was given as inputs in our risk-distributing practices. Put differently, while societies can make various choices about specific candidate behaviors for β and sufficiency thresholds for communication in various contexts, something like what I have outlined is needed as a social norm governing practices regarding consent.

The obvious question for this proposal is *what if the adopted social norms are bad?* What if, for example (as Hornsby and Langton and West (1999) argue), the prevailing conventions for sexual consent treat women's utterances of 'no' as meaning something like 'convince me', and take the absence of active, sustained and aggressive resistance to mean 'yes'? Would it follow on this view that we must, as Husak and Thomas (1992) suggest, consider it an empirical question whether women who said 'no' and meant *no* actually signalled consent, a question to be answered not by moral inquiry but by data about the statistical frequency of other women using 'no' to mean *maybe*?²⁵

While this may look like a weakness of my proposed view, it is actually its strength. Recall that to be able to communicate a moral content like consent, a signal must pass BURDENS: agents must be able to avoid the behavior without undue cost. If a convention treats *the absence of active resistance* as consent, the signal adopted is quite difficult (and indeed dangerous) to avoid, and so will only pass BURDENS if it's incredibly valuable to be able to communicate consent in that way, rather than in more easily avoidable ways. While we might have an interest like that in indicating consent to life-saving surgeries, it isn't plausible that the interest in indicating consent to a sexual encounter is strong enough to vindicate a signal that makes *consenting* the default. In general, the more stringent

communities. Exempting every solitary individual unfamiliar with local norms, by contrast, would only destabilize local conventions.

²⁵ Husak and Thomas argue that survey data (drawn from Muehlenhard and Hollabaugh 1988) suggests that the "women's consent-to-sex convention" involves saying 'no' while meaning to consent, and that as a result men who engage in sexual intercourse with a woman who explicitly refused their advances should be considered blameless *if* the practice is common enough that a man can rationally believe that a woman who says 'no' welcomes his advances. Husak holds an attitudinal view of consent, and takes the only question in such cases to be whether the man is blameworthy; if not, the law should take no interest in the case. As Archard (1997) points out, Husak and Thomas's argument involves (at least) two serious errors: (1) it moves uncritically from what the convention *is* to what errors we should excuse, and (2) it assumes that we can discover conventions by examining mere propensities.

the right being waived, the stronger agents' interests in avoiding inadvertently consenting, and so the more easily avoidable a signal must be to satisfy BURDENS. We can safely say, then, that in this case neither 'no' nor failure to aggressively resist can communicate consent.²⁶

None of that prevents a society from *in fact* relying on such a signal; it only means that the signal is unable to do the work it's been conscripted to do. It can't communicate consent, so those who rely on it are at best taking significant risks of, and at worst actively violating, the rights of non-consenting agents. Relying on a signal one *knows* to be insufficiently avoidable is thus blameworthy and has a high chance of resulting in impermissible activity, but what of signals that aren't *obviously* too costly? When only some of the agents in a society recognise that the conventions are unable to do the moral work, how should we proceed?

The question deserves some care. Judgments of responsibility for errors and miscommunications must be made with reference to the conventions actually adopted by the society—the fact that some behavior could have been used to signal consent is irrelevant to determinations of responsibility unless it is actually used as such a signal in the community. However, this does not mean that we should excuse all errors that are in line with actual conventions, let alone mere social regularities.²⁷ It is plausible that the perpetrators of errors arising from bad conventions are genuinely non-culpable: their cognitive habits have been shaped by the conventions; they have been taught by others in the community that the behavior α is a reliable signal of consent; they might be genuinely unaware that it fails BURDENS and is incapable of doing the necessary moral work. So, when they take α as evidence that X is permissible, they genuinely (though falsely) believe that they do not thereby transgress any rights.

A bad convention is problematic even if the errors are made non-culpably, and changing it requires destabilising the offending signals. That requires changing the reinforcement pressures by making the errors noticeably costly: agents who rely on α as signaling consent must face high costs when the other party did not in fact consent. This observation appears often in legal scholarship debating whether to move to a 'yes means yes' standard for consent. The worry raised in answer (which might concern you now) is whether it is fair or consistent with our practices of moral

²⁶ These considerations militate in favour of an affirmative consent standard for sexual encounters, insofar as the alternative treats mere absence of explicit refusal as constituting consent. Affirmative consent need not be limited to an explicit, verbal "yes"; California, for example, defines consent as 'positive cooperation in act or attitude pursuant to an exercise of free will' (4 CAL. PENAL CODE §261.6), which is an affirmative definition insofar as if A simply does nothing, it is not considered consent by default. For an illuminating discussion of the impact of affirmative consent standards in American criminal law, see Tuerkheimer (2016).

²⁷ A regularity is a pattern of behavior exemplified by a community, whereas a social norm or convention characterises the practices the population takes to be appropriate, even if they do not act accordingly. (Bicchieri (2006) characterises the latter as being composed of (a) empirical expectations and (b) a normative belief.)

accountability to hold non-culpable agents responsible for mistakes, requiring them to apologise to and compensate the agent who suffered the violation.²⁸

Let's start with whether holding B responsible for a mistake she made blamelessly could be fair. Given the informational limitations agents operate under, there will always be some non-culpable errors, and it can be tempting to think that B should be held responsible only if she was in some way culpable. That policy, however, leaves all costs of blameless error to be borne by the party whose rights have been violated. In holding A responsible for errors that result from his inadvertently communicating that he consents, we have already identified which costs of error we could reasonably demand that A bear. The remaining blameless errors are cases where it was merely rational for B to assume that A consented, and B was wrong. Mistakes resulting from bad conventions are a subset of these: the bad conventions skew B's evidence (or inference patterns) in a way for which B is not culpable, but which do not justify requiring A to simply bear the costs of the error. When the error propensities are shaped by bad social conventions, the problem compounds, as some agents are disproportionately likely to be the victim of non-culpable rights violations. From a distributive standpoint, to leave the costs on A renders him less secure against trespass than other agents, and the only justification we can offer for this is that members of the society are more prone to trespass against him than others. Insofar as A can reject this reasoning, it is fair to demand that those *making* the errors bear the additional costs. When B benefits from the error, and A doesn't, the case for holding B responsible for compensating A is even stronger.

Holding B responsible—and indeed, even blaming B—for errors resulting from bad conventions is also consistent with our moral practices of holding each other to account. As McGeer and Pettit (2015) argue, one of the functions of holding other agents accountable for moral failures is that in doing so, we maintain their sensitivity to reasons they can recognise, and put them into a position to respond to moral reasons to which they hadn't previously been sensitive. An agent who non-culpably relies on a bad signal is not yet in a position to recognise that such reliance is wrong; holding him responsible for the resulting error is one way of addressing him in an exhortation to do better *qua* moral agent.²⁹ Calling-to-account can be effectively used to alert agents to the fact that a convention they took to be unproblematic is in fact morally bankrupt. From this constructive vantage, it is appropriate to hold agents responsible for some non-culpable failures in order to put them and others in a position to recognise that the signals are bad ones.

²⁸ For example, Gruber (2016, 427) worries that holding agents responsible for mistakes that align with a bad, but active, convention amounts to suggesting that "we should randomly punish some who engage the script in the hope that it will change the world." For those concerned, let me note that I am not here advocating *criminalization* of rational errors. In other contexts, agents who negligently cause harm to others face civil, rather than criminal charges, aimed at compensating the harmed party rather than at punishing the offending agent. Something similar would seem appropriate in these cases, though it is not the focus of this article to give any specific legal policy recommendations. (See also Ferzan (2016))

²⁹ McGeer and Pettit (2015, 186)

4.4 Victim Blaming

The structure I've outlined here can also begin to explain at least some of the issues at stake in disputes over 'victim blaming', which occur when people respond to a victim's claim of having been wronged by noting that something about the victim's attire, history, or behavior explains the transgression — perhaps even makes it *reasonable*. Unanswered, discourse like this works to establish a social norm, exerting pressure on agents to rely on the cited behavior as a proxy signal for the relevant moral permission. If a concerned community member sees that a behavior cited in this way (e.g. wearing a short skirt as a signal of consent to sexual contact) fails the BURDENS constraint, then they should resist attempts by others to install it as a normative signal. Resistance requires refusing to accept it as a justification for errors, and refusing to make the performer bear the costs of errors that occur following the contested behavior; which requires insisting that the party who, however rationally, treated it as an indicator of consent, be held responsible. This is a high stakes project, since once socially accepted, signals shape agents' behavior: others will rely on the socially accepted signal to determine when they may act as though they have received consent.

My account might also clarify how to make progress in some of these disputes. It enables us to recognise that the cited considerations may diminish B's culpability, while rejecting the implication that they are reasons to deny that A was wronged, or is owed compensation. It allows us to see clearly that, of all the reasons which might lead B to mistakenly assume that A consents, only a small subset are of the kind that shift the costs of the moral gamble onto A. There are sometimes eminently rational gambles which we nevertheless lose, and it would be preposterous to argue that the fact that the gamble was rational means we have not in fact lost, and are not responsible to pay the stakes of the bet. Recognizing that the same holds for mistakes about consent puts us in position to shift the discussion to *which* considerations distinguish reasonable from rational errors, which is simply the project of identifying signals that satisfy the PUBLICITY and BURDENS constraints.

While I've urged that we must take seriously the fact that consent is communicative, we must also be careful in attaching moral consequences to extant communicative conventions, since socially operative conventions may be defective in a number of ways. We must, in short, accept that consent is communicative, but recognise that this gives us reason to examine and reform our communicative conventions.

5. Why Call this Consent

I mentioned at the outset that other theorists have seen that consent-based permissions must be observable, but were content to leave 'consent' as a name for agents' internal states, and call their accounts by other names. So why insist that my conclusions hold for *consent*, rather than following their example?

'Consent' names a particular moral power: the power to alter one's rights in a way that yields a consent-based permission. Whatever our analysis of the core cases, the normative interest grounding this power is only secured if we can reliably achieve uptake from other agents, which we cannot do

without addressing normative opacity. This is the fact that motivates a two-part analysis of consent of the kind I have outlined. Whatever moral story we tell to justify taking the core exercises to yield moral permissions—whether we ground the power of consent ultimately in agency, or autonomy, or something else—the very same reasons will justify taking the communicative behaviour to be sufficient to yield such a permission under the conditions I outlined. There is plenty of theoretical room for a mental state of ‘assent’, which requires that the individual competently decide to waive her moral complaint against or welcome a treatment X. But whether an individual is in such a mental state is not what determines whether they alter the rights-based permissions, prohibitions, or reasons of other agents. Rather than being identified with any purely descriptive thing like a particular internal state, or specific public performance, ‘consent’ names the moral power, and there are simply two ways to exercise this power: one which is identified by our core analysis (which may require internal assent), and one which is just a particular public behavior (communicating that one consents).

Taking ‘consent’ to be only the first of these leads us to divorce the receipt of consent-based permissions from what it takes to *give* consent. To preserve the coordination function of consent, we must allow that A’s performance of the observable consent-communicating behavior can be sufficient for B’s receipt of consent-based permissions. But if we then take a specific mental state to constitute or be necessary for consent, we have to say that moral consent-based permissions can be received even when moral consent is not given. My proposal allows us to *anchor* consent in intention—in all the core cases, when a person performs the consent-communicating behavior, it will veridically indicate their intending to consent—but also allows observable behavior to serve the coordination role. *Either* of these might be sufficient, on its own, for A to have morally consented, though only observable behavior shifts how third parties should distribute costs when risks eventuate. We thus escape having to say that there are any non-consensual consent-based permissions. The account has disjunctive implications, but has a unified underlying explanation of why each is a *way* of exercising the very same moral power. Rather than two classes of reasons, one justifying each disjunct, the very same considerations motivate both parts of the two-part analysis I have offered. We have a normative interest in being able to exercise our agency to intentionally give consent-based permissions also motivates taking communicative behaviour to be sufficient for genuinely changing permissions.

There is another, non-trivial reason to prefer my way of carving up the moral terrain. Alternatives which leave agents subject to considerable opacity about the boundaries of others’ rights start to push us towards a pernicious picture of how others’ rights should figure in agential deliberation, casting the goal as avoiding blame, rather than avoiding wrongdoing. In the grip of these pictures, we risk collapsing the question of whether A was wronged into whether B was blameworthy. The coordinating function of consent is to enable agents to avoid rights trespasses, not just to avoid culpability for trespassing. Facts about the moral permissions figure in agential deliberation in ways that facts about excuse cannot: that *A has consented* is a reason for B to consider

Xing permissible when deliberating; that *B will be excused* for Xing is not.³⁰ The facts that enter into B's deliberation about whether to X must be facts about permission, not excuse, and they can only be permissions if they track how A has changed his rights, not merely whether B is blameworthy or how society is disposed to evaluate B's action.

Though the view I have advocated gives up the assumption that consent must be intentional, it has many features which recommend it as an account of moral consent. It satisfies the core intuitive thoughts about what consent must involve (AGENCY and OBSERVABILITY). It fits into a unified analysis of the ways that moral powers operate, drawn from the considerations that justify the existence of moral powers in the first place. It explains why those powers appear to have a disjunctive structure, without having to posit a brute disjunctive power. It also makes available a much-needed degree of nuance for public discourse about mistakes in consent. And, in acknowledging the role of conventions, it makes explicit the moral constraints on normative conventions, thus lighting the way for evaluation and reform of our actual social practices. Most importantly, it provides the framing necessary to raise questions about what sorts of behaviors can be legitimately relied on to communicate consent, highlighting the broader social costs of communication when risks are asymmetric.

³⁰ Many authors stress this fact in support of communicative views; see for instance Dougherty (2015), Cowart (2004), Wertheimer (2003), Feinberg (1986b).

Acknowledgements

Thanks to Tom Dougherty, Jessica Keiser, Esther Klein, Massimo Renzo, and Jonathan Quong for comments that significantly improved this paper, and to Geoff Brennan, Daniel Halliday, Sukaina Hirji, Heidi Hurd, Karen Jones, Colin Klein, Ole Koksvik, Holly Lawford-Smith, Richard Rowland, Alex Sandgren, Mark Schroeder, Nic Southwood, and Daniel Wodak, as well as audiences at the 2018 Melbourne Applied Philosophy Seminar, the 2019 meeting of the Society for Philosophy, Politics, and Economics, and Virginia Tech University for helpful discussion of the ideas presented here.

References

- Alexander, Larry. 1996. "The Moral Magic of Consent (II)." *Legal Theory* 2 (2): 165–74.
- . 2014. "The Ontology of Consent." *Analytic Philosophy* 55 (1): 102–13.
- Archard, David. 1997. "A Nod's as Good as a Wink': Consent, Convention, and Reasonable Belief." *Legal Theory* 3: 273–90.
- . 1998. *Sexual Consent*. Boulder, Co: Westview Press.
- Bicchieri, Cristina. 2006. *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge University Press.
- Bolinger, Renée Jorgensen. 2018. "The Rational Impermissibility of (some) Racial Generalizations." *Synthese*
- Cowart, Monica. 2004. "Understanding Acts of Consent: Using Speech Act Theory to Help Resolve Moral Dilemmas and Legal Disputes." *Law and Philosophy* 23: 495–525.
- Dougherty, Tom. 2013. "Sex, Lies and Consent." *Ethics* 123: 717–44.
- . 2015. "Yes Means Yes: Consent as Communication." *Philosophy and Public Affairs* 43 (3): 224–53.
- . 2018a. "Affirmative Consent and Due Diligence." *Philosophy and Public Affairs* 46: 90–112.
- . 2018b. "On Wrongs and Crimes: Does Consent Require Only an Attempt to Communicate?" *Criminal Law and Philosophy*.
- Feinberg, Joel. 1986a. *The Moral Limits of the Criminal Law: Harm to Self*. Vol. 3. Oxford University Press.
- . 1986b. "Victim's Excuses: The Case of Fraudulently Procured Consent." *Ethics* 96 (2): 330–45.
- Ferzan, Kimberly Kessler. 2016. "Consent, Culpability, and the Law of Rape." *Ohio State Journal of Criminal Law* 13: 397.
- Fuller, Lon L. 1969. *The Morality of Law*. Revised. New Haven: Yale University Press.
- Gruber, Aya. 2016. "Consent Confusion." *Cardozo Law Review* 38: 415.
- Guererro, Alex. 2019. "The Epistemology of Consent." In *Applied Epistemology*. Edited by Jennifer Lackey. Oxford University Press.
- Hartogh, G den. 2011. "Can Consent Be Presumed?" *Journal of Applied Philosophy* 28 (3): 295–307.
- Hickman, S. E., and C. L. Muehlenhard. 1999. "'By the Semi-Mystical Appearance of a Condom': How Young Women and Men Communicate Consent in Heterosexual Situations." *The Journal of Sex Research* 36: 258–72.

- Hurd, Heidi. 1996. "The Moral Magic of Consent (I)." *Legal Theory* 2 (2): 121–46.
- Husak, Douglas. 2006. "The Complete Guide to Consent to Sex: Alan Wertheimer's "Consent to Sexual Relations"." *Law and Philosophy* 25 (2): 267–87.
- Husak, Douglas, and George Thomas. 1992. "Date Rape, Social Convention, and Reasonable Mistakes." *Law and Philosophy* 11: 95–126.
- Langton, Rae, and Caroline West. 1999. "Scorekeeping in a Pornographic Language Game." *Australasian Journal of Philosophy* 77 (3): 303–19.
- Liberto, Hallie. 2017. "Intention and Sexual Consent." *Philosophical Explorations* 20: 127–41.
- Manson, N.C. 2016. "Permissive Consent: A Robust Reasons-Changing Account." *Philosophical Studies* 173: 3317–34.
- McGeer, Victoria, and Philip Pettit. 2015. "The Hard Problem of Responsibility." In *Oxford Studies in Agency and Responsibility*, edited by David Shoemaker, 3:160–88. Oxford: Oxford University Press.
- Miller, Franklin, and Alan Wertheimer. 2009. "Preface to a Theory of Consent Transactions: Beyond Valid Consent." In *The Ethics of Consent: Theory and Practice*, edited by Franklin Miller and Alan Wertheimer, 79–105. Oxford University Press.
- Muehlenhard, Charles, and Lisa Hollabaugh. 1988. "Do Women Sometimes Say No When They Mean Yes? The Prevalence and Correlates of Women's Token Resistance to Sex." *Journal of Personality and Social Psychology* 54: 872–79.
- Owens, David. 2011. "The Possibility of Consent." *Ratio* 24 (4): 402–21.
- . 2012. *Shaping the Normative Landscape*. Oxford: Oxford University Press.
- Sherwin, Emily. 1996. "Infelicitous Sex." *Legal Theory* 2: 209–31.
- Shiffrin, Seana. 2008. "Promising, Intimate Relationships, and Conventionalism." *The Philosophical Review* 117 (4): 481–524.
- Simmons, A. John. 1981. *Moral Principles and Political Obligations*. Princeton University Press.
- Skyrms, Brian. 2010. *Signals*. Oxford University Press.
- Tadros, Victor. 2016. "Causation, Culpability, and Liability." In *The Ethics of Self-Defense*, edited by Christian Coons and Michael Weber, 110–30. Oxford University Press.
- Tuerkheimer, Deborah. 2016. "Affirmative Consent." *Ohio State Journal of Criminal Law* 13: 15–49.
- Watson, Gary. 2009. "Promises, Reasons, and Normative Powers." In *Reasons for Action*, edited by David Sobel and Steven Wall. Cambridge University Press.
- Wertheimer, Alan. 2003. *Consent to Sexual Relations*. Cambridge University Press.
- Westen, Peter. 2004. *The Logic of Consent: The Diversity and Deceptiveness of Consent as a Defense to Criminal Conduct*. Aldershot: Ashgate Publishing Company.