

Objective Rights & Epistemic Risks: Beyond Subjective Permissibility

Renée Jorgensen *

1 Ignorance, Rights, and Right Action

At 2 am, four plainclothes NYPD officers spotted a car matching the description of the vehicle of a serial rapist, believed to be armed and dangerous. A man was on the porch of a nearby residence, strongly backlit. They approached, commanding him to show ID. He reached into his jacket, and several things happened in quick succession: the closest officer yelled 'gun!', there was a loud bang, and that officer fell forward on the porch stairs. The remaining officers fired in response, killing the resident.

Given their epistemic position, the officers rationally thought that by firing, they were imposing necessary and proportionate **justified** defensive harm (ergo not violating the resident's rights) to protect the fallen officer. But they were mistaken. In fact, the resident posed no threat at all: he was only retrieving his wallet— not a gun—and the fallen officer had simply tripped thunderously on the stairs.¹ So the officers' decision to fire was an **error**: an action done in non-culpable ignorance, that, if performed with full knowledge of the situation and consequences, would be a rights-violation. Cases like these occur at semi-regular intervals; when they do, some argue that since a real gunman who was bluffing or a bad shot (and therefore posed no genuine threat) would not be wronged by a similar error, the facts determining whether an error violates rights must be **what evidence is available to the duty-bound**, rather than what the rights-holder actually did, and therefore the police action actually did not wrong the resident. This is a condensed line of reasoning, possibly with more support from sophisticated moral theory than one might expect, and well-worth interrogating. And even though most of us do not regularly face decisions about whether to use lethal force, we do usually operate with imperfect information and limited cognitive power; we routinely lack knowledge necessary for avoiding error in the many lower-stakes decisions we make. So it is important to ask whether our non-culpable ignorance matters, and what (if any) difference it makes to what we owe each other.

We should distinguish this question—whether or how ignorance changes what an agent is duty-bound to do **vis a vis** others' objective moral rights—from a superficially similar, but very different question: whether ignorance changes what they are morally permitted or required to do all-things-considered. To respect Abe's moral rights, Bea must refrain from activities that Abe has a right against (e.g. harming him, taking his property, etc.), and must perform actions to which Abe has a positive right (e.g. fulfilling her promises to him, giving adequate deliberative weight to Abe's moral claims). We can think of the set of these demands as

* I am indebted to Alan Hájek, Rahul Kumar, Daniel Muñoz, and Shmulik Nili, Anni Raty, Chelsea Rosenthal, Holly Smith, and Jonathan Quong for helpful written comments on earlier versions of this paper. I am grateful as well to discussions with Joe Bowen, Stephanie Collins, Steve Darwall, Julia Driver, Bob Goodin, Tom Hurka, Frank Jackson, Frances Kamm, Julia Markovitz, Tom Parr, Philip Pettit, Mark Schroeder, Kieran Setiya, Nic Southwood, and Alex Worsnip, as well as the participants at the 2018 workshop on Ethics & Risk at Australian National University, the 2019 MIT workshop on Moral Rights, and 2022 symposium on Epistemic Constraints on Moral Theory at the Eastern APA, which helped to refine the ideas presented here. I've presented the early ancestors of this paper at Otago University, the

University of Melbourne, National University of Singapore, and Wuhan University, and thank those philosophers for spurring me to develop these ideas. This draft was last updated March 15, 2022. Comments are welcome (rjorgen@umich.edu), but please do not cite or circulate without permission.

¹This case is based on the police report for the NYPD's shooting of Amadou Diallo, which occurred on February 4, 1999, but I have simplified it to bring out the moral questions more starkly.

1

Objective Rights & Epistemic Risks²

constituting a rights-based standard that **allows** an action iff it does not breach any duty correlative to those rights. Such a standard is narrow in scope and coarse-grained in prescription: all it does is identify whether an action is allowed or forbidden by the rights-based reasons. It says nothing about other sources of value or what would be the best thing to do. By contrast, a standard of **right action** reflects what morality writ large requires of an agent: it takes inputs of all kinds, and outputs what the agent morally ought or is permitted to do, all things considered. It may also yield fine-grained prescriptions, ranking an agent's options from best to worst according to their comparative moral value, and obliging her to do the best thing she can.² And while a standard of right-action identifies a monadic property of actions (viz., **being obligatory**), the duties reflected by a rights-based standard have a dyadic structure; they are actions **the agent owes to a particular rights-holder**. Accordingly when an agent violates the dictates of a rights-based standard, she violates some particular person's moral right, **wronging them**.³ When she violates the dictates of the standard of right action, she simply **does something morally wrong**. Because of the differences in their scope and operation, the demands of the two standards can diverge. Many things might be morally impermissible which do not violate anyone's rights, and some rights-violations may be all-things-considered permitted or even required.⁴ To clearly distinguish the two, in what follows I will say that an action is 'allowable' if it does not violate or infringe a right, reserving 'permissible' for all-things-considered evaluations.

So put more carefully, our question is **when, if ever, does non-culpable agent-ignorance change whether an action is allowable?** Moral theorists often assume an **Outcome-Relative** picture in answer: We determine whether B's action violates A's rights (wronging him) by asking whether its outcome is one that A had a weighty moral interest in escaping. If B were to knowingly bring about that outcome, would her action clearly violate a duty she owes to A, wronging A? If so, then B's acting in a way that in fact causes the outcome violates A's rights. Whether B **actually** could have guessed that her action would have these results matters for whether we blame her, but is irrelevant to the question of whether A was wronged. On this picture, the rights-based standard demands that the agent do what she **would have a duty to do, were she to know all the relevant facts**.⁵ Her actual ignorance makes it harder to fulfill her duties, but it does not change their contents; respecting others' moral rights requires avoiding errors, **simpliciter**. So while non-culpable ignorance may excuse a mistaken agent from blame, it never changes whether an action is allowable. According to this view, because the

officers' decision to fire killed the resident when he was not in fact liable to lethal defensive harm, their error violated his rights— and the same would hold for a non-liable bluffer or bad marksman.

Some critics of this picture object that such insensitivity to epistemic limitations holds the duty-bound agent to an impossibly high standard. They suggest that rather than having rights to certain **outcomes**, rights

²This characterization is meant to remain neutral between a dizzying array of consequentializing and deontological standards of right action; they are not our focus.

³For simplicity I am setting aside the possibility, floated by [Cornell \(2015\)](#), that wrongings come apart from rights in certain cases.

⁴When lesser-evil considerations justify violating an agent's rights, the standard of right-action requires something that the rights based standard prohibits; violations justified in this way are typically called 'infringements'. Those (like [Oberdiek, 2004](#)) who reject the notion of permissible rights-infringements will hold instead that rights are specified in such a way that they do not prohibit (and thus are consistent with) the actions that other theorists take to be justified infringements.

⁵This view is discussed by other scholars under the label 'fact-relative' or 'Objective View' (See, e.g., [Oberdiek, 2017](#); [Quong, 2015](#); [Thomson, 1990](#); [Zimmerman, 2008](#)), but these labels invite confusion by implying that a fact-relative or objective standard must necessarily treat agent-ignorance as irrelevant. So I prefer to refer to this as **outcome-relative** view.

Objective Rights & Epistemic Risks 3

holders are entitled to others' **deliberate conduct**. A's moral interest in escaping an outcome gives B a duty to do her best to avoid causing it, given the evidence available to her. On this view, what the rights-based standard demands from agents is determined **relative to their evidence** (or the evidence they would have, were they appropriately attentive and rational). So non-culpable ignorance always affects whether the action is allowable. According to this view, because the officers' evidence (misleadingly) implied that firing would fulfill a duty they owed to the fallen officer without violating a duty owed to the resident, their error did not violate his rights, though it unfortunately killed him.

One might reasonably think that the truth lies somewhere between 'always' and 'never'; that we can affirm that the resident's rights **were** violated, but **also** say that agents' non-culpable ignorance affects whether (for instance) a bluffer or bad marksman have grounds to complain. Proposals like this are reasonable, and indeed quite a few have been made—but are quickly distorted by an assumption that the two pictures exhaust the options. Criticisms of the outcome-relative picture are presented or interpreted as objections to **any** fact-relative standard for rights; proposals that agents' epistemic limits are **sometimes** relevant are quickly glossed as advocating an **evidence-relative** picture.⁶ I suspect that this is partly due to the influence of the adjacent debate over whether the standards for right-action are set by an agent's 'subjective circumstances'— her beliefs, evidence, etc., as Subjectivists contend—or her 'objective circumstances'—all the facts except those concerning her beliefs, evidence, etc., as Objectivists maintain. But while these two discussions have deep similarities, it is treacherous to assume that all of the considerations raised and conclusions established for right-action will carry over to the debate about rights. My goals in this paper are principally negative; I aim to dislodge three common assumptions, which I take to both fuel the condensed

argument and distort discussion:

1. that the outcome-relative and evidence-relative pictures exhaust the options for theorizing about rights, such that rejecting one forces us to accept the other; relatedly,
2. that taking agents' rights to be fact-relative entails that the epistemic limitations of the duty-bound speak only to blame, and not to whether her action is allowable; and
3. that common defenses of Objectivism about right-action can be readily deployed to defend outcome relativism about rights against parallel objections.

I'll briefly recount the outcome- and evidence-relative characterizations of rights in §2, highlighting the inheritances from debates about right-action. Then I will argue that the inherited framework doesn't fit. The distinctive structure and logic of rights-based obligation (surveyed in §3) simultaneously impose more constraints on and introduce new options for theorizing. I'll demonstrate (in §4) that familiar as it may be to take 'fact-relative' to mean 'outcome-relative'—and whatever the merits of the parallel move might be for right-action—it is a mistake when theorizing about rights. For objective moral rights to serve their core theoretical roles, they must be both fact-relative and at least somewhat sensitive to agents' epistemic limits.

2 Facts and Evidence

The thought that an agent's ignorance cannot affect what an objective standard demands from her has too long a history to recount in any real detail. It has principally been defended for standards of right-action;

⁶See Frowe (2015); Quong (2015); van der Vossen (2016); Zimmerman (2006) for examples of the former; Oberdiek (2017); T. Scanlon (2008) for the latter.

Objective Rights & Epistemic Risks⁴

what morality writ large requires of an agent. G.E. Moore (1903) famously embraced this view, writing that “. . . the term **duty** is certainly so used that, if we are subsequently persuaded that any possible action would have produced more good than the one we adopted, we admit that we failed to do our duty.”⁷ Ross (1930 [2002], **The Right and the Good**, p.32) extended the outcome-relative line to include duties arising from rights, holding that what such duties demand is “that which if I were omniscient I should see to be my duty.” And even though Ross repudiated this position to defend an evidence-relative view a few years later (in **Foundations of Ethics** Ross, 1939), this appeal to the omniscient perspective lives on in the contemporary heuristic representing fact-relative moral requirements as aligning with the advice one might receive from a morally motivated, **fully-informed** adviser (deployed by many, including Thomson (1990) and Parfit (2011, 151-2)).

The terrain here gets slippery, and theorists slide quickly from the fact that a standard or duty is 'objective' or 'fact-relative' to the conclusion that it is sensitive only to the facts about what the actual outcomes are or would be. It is even common to take complete insensitivity to agents' epistemic states to be the **signature** feature of objective moral standards. Hurd (1999, p. 1571), for example, writes that

“justifications concern themselves with the rightness of actions, which is an 'objective' determination in the sense that it proceeds without reference to the subjective mental states of actors; and excuses concern themselves with the culpability of actors, which is a 'subjective'

inquiry in the sense that it is interested in actors' mental states and unique epistemic limitations."

Similarly McMahan (2009, p. 43) presents as a definitional, rather than substantive, the claim that "An act is objectively permissible when what explains its permissibility or justifiability are facts that are independent of the agent's beliefs." And Graham (2010, p 88) defines an agent's objective situation as consisting "of all the facts about her at that time except those about what evidence she has about the world."

Theorists under the influence of this picture treat fact-relative rights as **by definition** insensitive to agent ignorance. For instance, Tadros (2016, p. 112) maintains that a reasonably mistaken agent "Y wrongs X in the fact-relative sense as Y does what he ought not to do were he to know all the relevant facts." Thomson (1990, p. 233) explicitly reasons from the fact that agents "will infringe a claim if they know they will cause these outcomes to come about, and nevertheless proceed anyway" to the conclusion that doing so in non culpable ignorance is equally an infringement.⁸ And, though he ultimately argues against the view, Quong (2015, p. 249) glosses fact-relative accounts of rights as constitutively committed to the view that

"Whether B infringes A's right not to be harmed by ϕ -ing depends on what will in fact occur if B ϕ s. B's knowledge of, or evidence about, the exact consequences of her ϕ ing are irrelevant to the question of whether her ϕ -ing constitutes an infringement of A's right not to be harmed by B."

Though far from comprehensive, this is a representative sample demonstrating how commonly it is assumed that to determine what a fact-relative moral standard requires of an agent, we must ask a counterfactual: what **would** it require from her, **if she knew all the relevant facts?**⁹

⁷Moore (1903, *Principia Ethica*, 149-50)

⁸Thomson (1990, *Realm of Rights*, Chapter 9 ('Harm'), especially pp. 233-245).

⁹There are some important exceptions; Ferzan (2005) and Frowe (2010) argue that epistemic facts are relevant to objective defensive rights, but endorse the omniscient interpretation for other rights.

Objective Rights & Epistemic Risks 5

This encodes an outcome-relative characterization of rights like the one Thomson (1986) advocates: your rights against harm put me under a duty **to not cause you to be harmed**. Thus any action by which I **in fact** cause harm to you infringes your right—even if no one could have guessed **ex ante** that harming you was a possible result of my action. Similarly, even if an action of mine posed high risk of harming you, and even if I undertake it with the intention to harm you, it does not infringe your right unless it actually causes the harmful outcome.¹⁰ The outcome-relative view is tidy: it cleanly separates the facts affecting what is allowable, on the one hand, from the facts about agents' epistemic situations, on the other.

The evidence-relative view, by contrast, takes agents' epistemic positions to completely determine objective rights. Zimmerman (2008) outlines a Prospective View for both moral right-action and rights-based obligations, on which an action is morally obligatory in general if

and only if it “provides the agent with a prospect of doing what suits his situation as a whole better than that provided by any other option he has.”¹¹ This approach begins by treating agents’ epistemic limitations as decisive constraints on their duties, and reasons backward to characterize rights as the corollaries of these constrained duties. So understood, rights against being harmed are the corollaries of our duties to not put others at unjustified risk of grave harm, given our evidence. When our best evidence suggests that harm is justified, we do not violate these duties if we impose harm. Similarly, if our on evidence an action ϕ is highly unlikely to cause harm, then we do not violate our duty by ϕ ing, even if unfortunately it does actually harm someone. In short, on an evidence relative view “one’s right to life, and, indeed, one’s rights in general are hostage not only to the abilities but also to the evidence possessed by those against whom one’s rights are held” (Zimmerman, 2008, p.112).¹² This is also a tidy view: B’s duties are a function of her evidence, and A’s rights are whatever correlate with B’s duties.

These views echo objectivism and subjectivism, respectively, about right-action. And many of the arguments rallied in favor of each (or against its opponent) are borrowed from the adjacent debate. But moral rights have a distinctive structure and function, so to evaluate the success of any characterization of rights we need to look at how well it fills the functional role.

3 Theoretical roles of rights

The substantive agreement about what moral rights are and how they operate can be organized into four or five core roles. Moral rights are grounded in agents’ moral status: people are owed certain consideration in virtue of their standing as agents. Without committing to any particular theory of rights, we can say that their **function** is to give agents “important forms of protection and control over their own lives” (T. M. Scanlon, 2003, p. 4). The sorts of rights relevant to our inquiry are those that restrict others’ actions; in Hohfeldian terms, these are **claim rights**.¹³ To say that an agent has a right in this sense is to say that there is “sufficient

¹⁰Thomson is thoroughly consistent on this point, denying that rights-claims require any expression of respect from the duty-bound, or must play any role in shaping their deliberation. Not all of her followers have been equally careful.

¹¹Zimmerman (2008, p. 78).

¹²van der Vossen (2016) focuses narrowly on defensive rights, but pursues a roughly parallel line, contending that when ‘B is liable’ is part of A’s “morally best bet” given her evidence, A is allowed to impose defensive harm on B. And if defensive harm is allowed, then A was not under a duty not to impose it, and therefore B had no right against suffering it.

¹³Hohfeld (1919) regiments rights into four types. Two are first-order rights: A has a **privilege** to ϕ if A is not under a duty to refrain from ϕ ing; she has a **claim** to ϕ if others have a duty either to not interfere or to enable her ϕ ing. Two are second-order: she has a

Objective Rights & Epistemic Risks 6

reason for holding some other person(s) to be under a duty.”¹⁴ This is because a claim right implies a correlative duty, a demand the rights-holder can morally make on others, and which she can reproach them for failing to fulfill. These rights provide the holder freedom and security by safeguarding her control over certain aspects of her life, ensuring that she won’t be deprived of important choices or goods by other agents, and guaranteeing that if she is, she has claim to compensation. Their value is thus partially constituted by the ways they induce others to regulate their conduct concerning her. This gives us a first theoretical role:

Security – objective rights function to secure the relevant moral goods for rights-holders.

Many of the most important moral goods protected by rights — security, respect, and the conditions of autonomy — can only be enjoyed if the rights-holder is not only presently not suffering a violation, but has a reasonable expectation against suffering future rights-violations. This is part of what justifies agents' permissions to impose proportionate costs to prevent would-be violations.

Duties implied by claim-rights are (typically) **enforceable**, meaning that it is allowable to impose certain costs on the duty-bound to prevent or compensate for a breach of their obligations. As [Feinberg \(1966, p. 143\)](#) puts it, “dutiful actions and omissions called for by the rights of other people [. . .] can be demanded, claimed, insisted upon, without embarrassment or shame. When not forthcoming, the appropriate reaction is indignation”. Similarly, [Darwall\(2006, p.18\)](#) writes that a claim-right “includes a second-personal authority to resist, complain, remonstrate, and perhaps use coercive measures of other kinds, including, perhaps, to gain compensation if the right is violated.” It is because of this connection to reproach and enforceable duties that **what rights there are** is importantly constrained by costs to the duty-bound: whether A has a right to B ϕ ing can depend on how difficult ϕ ing is for B.¹⁵

Correlativity – objective rights correlate with the objective duties others can be reasonably demanded to perform.

The connection is particularly vivid in debates over rights of rescue. Because rescuing someone can be costly (it at least takes time and effort, and may be hazardous), some theorists deny the existence of rights of rescue altogether. Others offer a balancing approach: if rescuing A from dire straits would be easy for B, A has a right to be rescued, but if it would be quite costly, rescuing A would be supererogatory for B; A does not have a right to it.

Rights also have an important distributive role: they represent the fair demands agents can make on one another. When what it is fair to demand depends on the background conditions, the contours of the corresponding rights-claims also depend on these conditions. This is nicely illustrated in [Nozick \(1974\)](#)'s account of property rights. He denies that unequal distributions are unfair, so the fact that private property rights yield an **inegalitarian** distribution of holdings is no objection to them. But it remains essential that they be **Fair**, which he defines as satisfying a Lockean proviso that acquisition must not worsen the position of others. While Nozick's characterization of property rights under moderately ideal conditions includes the right to exclude others or 'charge what he will' for access to one's property, if under the actual circumstances

power with respect to ϕ if she has the authority to alter others' first-order rights regarding ϕ , and she has an **immunity** with respect to ϕ if no one else has a power to alter her first-order rights with respect to ϕ .

¹⁴[Raz \(1986\)](#), **The Morality of Freedom** (Oxford: Oxford University Press), p. 166.

¹⁵Relative to a moralized baseline. As [Quong \(2015\)](#) rightly points out, the fact that A will die if he does not steal B's internal organs makes it costly for A to refrain, but does not in itself release A from an obligation to refrain from stealing B's organs.

Objective Rights & Epistemic Risks 7

these exclusions would violate the proviso, the contours of legitimate property rights change

accordingly.¹⁶ He is explicit that it's not that agents have more extensive rights which are simply outweighed or may be permissibly infringed; whether an agent **has** a right with certain contents depends on the background facts being a certain way. Whatever objective rights may require, they do not make morally unfair distributive demands. In short,

Distributive Fairness – the demands of objective rights do not result in an unfair distribution of the moral goods protected.

Fourth, once we know how various risks have developed, claim-rights partially determine who should receive, and who owes, compensation for costs incurred. Agents whose rights were infringed have a special moral claim to compensation or restitution that others lack.

compensation– Objective rights identify which agents have special claim to compensation or remedy **ex post**.

These four are features of the conceptual profile common to all rights-theorists. Nearly everyone also takes rights to play a crucial deliberative role, though some (notably Thomson, 1986) eschew this additional commitment. While an agent's rights may exist to safeguard her interests in certain states of affairs, she is not directly entitled to that **state of affairs** obtaining.¹⁷ Rights are fundamentally dyadic; they are claims to certain consideration and behavior **from agents**. One's interests can be injured without rights-violation: Abe's right against being harmed is not violated if he is mauled by a wild tiger, nor is his property right violated if a wildfire destroys his house (though it would be if an arsonist set the fire). Conversely, if an agent's deliberation manifests a disregard for the lives, property, or well-being of others—say by deciding to play Russian Roulette on unwilling participants for amusement—then even before the outcome is set, he has already failed to respect their rights, wronging them. To respect another's rights against being harmed it is not sufficient to merely **in fact** not harm her; one must treat the fact that she has a right against being harmed as a reason not to act in a way that would probably harm her.¹⁸ This yields a final constraint,

Deliberative Guidance – rights provide reasons for action that agents can appreciate and use for deliberation before knowing what the specific downstream consequences of their actions will be.

This function is more explicit in some formulations than others, but is clearest when rights are cast as imposing side-constraints on the means one can take to accomplish various ends, or as exclusionary or 'trumping'

¹⁶Nozick (**Anarchy, State, and Utopia** 1974, p. 180) clarifies this by considering what would happen if an agent "possesses one [water hole], and unfortunately it happens that all the water holes in the desert dry up, except for his. This unfortunate circumstance, admittedly no fault of his, brings into operation the Lockean proviso and limits his property rights. . . . Notice that the theory does not say that owners do have these rights, but that the rights are overridden to avoid some catastrophe. (Overridden rights do not disappear; they leave a trace of a sort absent in the cases under discussion.) There is no such external (and **ad hoc?**) overriding."

¹⁷Adherents of the Choice Theory may construe 'interest' here narrowly, as picking out only an agent's autonomy interests. Those partial to the Interest Theory can read the term more broadly; nothing in the argument to follow depends on either interpretation.

¹⁸Many theorists also accept the converse: if an agent exercises adequate caution but his action **unforeseeably** harms B, he does not infringe B's rights. McMahan (2005), for instance, restricts what counts as a rights-violation to foreseeable consequences of actions: if it is unforeseeable that the agent's ϕ ing will cause a given harm to some other person A, that harm is disvaluable but it is not a rights violation. His reason for this

is that unforeseeable consequences of ϕ ing are not something for which A is agentially responsible, and so should be treated like harms from tigers and wildfires.

Objective Rights & Epistemic Risks 8

reasons shaping which options other agents may treat as live when deliberating about what to do.¹⁹ There is a deep tension between this and the outcome-relative account of rights. Thomson herself escapes it by denying that any of the features of deliberative guidance are features of objective rights, but it is a significant challenge for theorists who aim to extend her general account without making similar sacrifices.

4 Measuring up?

Allow me to make a quick structural point about these roles before going further. security and distributive fairness pull toward a fact-relative approach, since there is reason to worry that an evidence-relative interpretation will render agents' entitlements too fragile to genuinely protect their moral interests. compensation seems similarly to track the interests of the agent, rather than the decision context of the duty-bound. But correlativity and deliberative guidance pull the other way: if rights are to imply reasonably demandable duties and provide **ex ante** reasons for action to the duty-bound, it seems likely they will need to be sensitive to agents' epistemic limits. This should make us skeptical that either taking agent-ignorance to **always** or **never** affect what is allowable will fill all five roles; we should expect to need a middle way.

4.1 Security & Distributive Fairness

Evidence-relative views struggle mightily with security and compensation. There are plenty of cases where non-culpable ignorance does not intuitively change whether an action infringes a right. If Bea mistakes Abe's coat for her own and takes it home, Bea **in fact** deprives Abe of his property without authorization, and surely must remedy this by returning the coat; Bea's belief, even if rationally supported by her evidence, does not change Abe's property rights. Some have suggested to me that the reason for this intuition is that nothing of much moral significance is at stake—it's just a coat, after all.²⁰ But similar intuitive verdicts arise with much higher-stakes. Recall the opening case: the police officers' best evidence justified their thinking that the resident was liable to—and thus lacked a right against—being defensively killed. Nevertheless, surely he **was** wronged when killed; he had in fact done nothing to incur a liability, and was even attempting to cooperate. An evidence-relative theory of rights cannot vindicate this intuition for the case as I have described it. Given the officers' evidence, it entails that the resident objectively lacks rights against, and **thus is neither wronged by, nor entitled to compensation or recompense for** being killed by the police.²¹

But before embracing outcome-relativism, we should take a closer look at how it actually fares with re

¹⁹For the former, see i.e. Nozick (1974), T. Scanlon (1998), Ripstein (1999), Darwall (2006), Tadros(2011), Cornell (2015); for the latter, Raz (1986), Kramer (2001), Dworkin (1978).

²⁰Frank Jackson made this suggestion in personal correspondence.

²¹Quong (2012, 61-2) makes a similar critique of the evidence-relative view by pointing out that evidence-relative views imply that the civilians are liable to be killed in a case he calls **Duped Soldiers**. A nuanced evidence-relative

standard can permit the resident to counter-defend without violating the officers' rights, but that is the most it can offer. To secure counter-defensive permissions, van der Vossen (2016, p.142) gives a disjunctive condition for liability: "Y being liable to defensive force is part of X's morally best bet if X is justified in believing either that Y is about to perform an objectively rights-violating act, or Y has culpably acted in a way that would, if successful, be such a rights-violating act, or both." Since the officers are justified in believing that the resident is liable, on this account the resident is liable to —lacks a right against— their shooting. So we cannot understand 'objectively rights-violating act' as an act that would **actually** violate a right. Instead, van der Vossen clarifies that these are actions that "would violate someone's rights, and thus be something that Y (other things equal) ought not to do, were Y to have all the relevant evidence available." In other words, errors. The account thus can ground the claim that the resident is permitted to counter-defend, but only by attaching liabilities to counterfactual rights violations: the officers do not actually have full-information, so the resident does not actually have a right that will be violated by their action (assuming an evidence-relative rights-based standard).

Objective Rights & Epistemic Risks⁹

spect to characterizing agents' fair demands on each other. To make this determination, we must examine not only what it directly represents agents as entitled to, but also what it actually secures for them: how it distributes the risks of suffering (intentional or ignorant) violations under their actual circumstances. This depends in part on what deliberative guidance it offers. An advocate of outcome-relativism has two options here. She can either austere maintain that agents should follow the objective standard directly, avoiding all errors regardless of their evidential (im)probabilities, or hold that agents should be guided by it only in directly by relying on a standard of subjective permissibility that instructs them to discount risks of error by their subjective (or evidential) probabilities. Either option expectably leads to an unjust distribution of the moral goods protected by rights.

4.1.1 Austere Guidance

Let's start with Austere Guidance. Rights against harm are secured by a privilege to impose necessary and proportionate harm to avert an unjust threat. Agents are often uncertain whether an apparent threat is genuine, and hence uncertain whether imposing defensive harm would be an allowable exercise of their own rights, or a violation of someone else's right against harm. Imagine an agent deciding whether to defend against what is very probably an unjust threat:

	.9	.1
	genuine threat	no threat
Defend	permissible	defense error
Refrain	suffer harm	no harm

Were she weighting her alternatives by evidential probability, this agent should defend herself. But on the present suggestion, the agent avoids violating A's rights if and only if she **actually avoids an error**— she wrongs A precisely as gravely in making an unlikely error as in making a certain one. Given that, if there is an option that allows her to eliminate the risk of violating A's rights, then a conscientious agent must choose it unless there is at equally morally weighty counter-consideration. The instruction to disregard evidential probabilities constrains her deliberation to comparing the **ex post** values (outcomes) of each alternative.²² In this case, the best of the good outcomes is <no harm>, in which the agent refrains and the other party is no threat. The best bad outcome is <suffer harm>, because <error> constitutes violating a rights, which, all else equal, is morally worse than suffering a similar violation.²³ Choosing the action with the best good outcome, and least bad worst outcome in this way will lead an agent to refrain from enforcing their rights against harm

whenever they are uncertain—and they will nearly always be uncertain. So while this theoretically accords agents rights against harm, they lack security, since it practically always denies them permission to enforce these rights.²⁴

²²To those familiar with the debate, it will be clear that this interpretation of the requirements of rights is parallel the Deontic Moral Absolutism criticized by Jackson and Smith (2006).

²³This follows from even the weakest form of the doing/allowing asymmetry. On classic articulations, the fact that it would save five innocent lives is not sufficient to justify killing one innocent person, so the moral disvalue of killing must be several times as difficult to justify as the disvalue of the harm of allowing to die. (See discussions in Foot, 1967; Quinn, 1989; T. Scanlon, 2008, p.91).

²⁴Defensive permissions are not uniquely troublesome; similar problems arise for when agents face uncertainty concerning whether they have received consent, or been commanded, or made a promise.

Objective Rights & Epistemic Risks 10

4.1.2 Subjective Guidance

This problem is sufficiently obvious that few (if any) advocate Austere Guidance. The preferred alternative is a hybrid approach, which still holds that whether an action is a rights-violation is set by an objective, outcome-relative standard, but directs agents to deliberate according to a subjective evidence-relative standard. The effect is to instruct agents to take rational risks: acting when the evidential probability of error or the expected moral disvalue is sufficiently low.²⁵ This aligns closely with the guidance an evidence-relative view gives to agents. The difference is that the evidence-relative view takes actions that are subjectively permissible to be objectively allowable, whereas the hybrid outcome-relativist is instructing agents to **aim to avoid culpability**, rather than to avoid objective wrongdoing. Unfortunately both views will expectably yield an unfair distribution of moral risk (violating the Distributive Fairness constraint), but for very different reasons from what we've just rehearsed for the austere guidance approach.

To see how, focus on how well-intentioned but ignorant agents make errors. Our evidence includes a wide variety of facts; among them, the many things we've learned about how people in various social groups tend to behave. Background beliefs of this kind can strongly influence an agent's rational credence about whether A is acting in a way that makes it allowable for her to ϕ , particularly if A is a member of an especially visible social group. Suppose, not improbably, that members of the working-class are over-represented among perpetrators of mugging theft, and Bea knows this. Seeing that Abe is working-class then shapes her baseline confidence about whether Abe intends to rob her: she'll assign it higher probability than she otherwise would. As a result, she'll require less confirming evidence before acting on the assumption that he is a threat. More insidiously, she'll also be disposed to interpret ambiguous information—like a sudden loud noise or rectangular shape—through this expectation, and so sincerely believe she heard a gunshot or saw a gun (just as happened in the opening case). Relying on their rational credences leaves even the best intentioned agents disproportionately likely to make errors when interacting with members of certain salient social groups; to misread men from racial minorities as aggressive threats, or women as consenting.²⁶ Under these conditions, members of these groups face disproportionately high risk of suffering rights-violations not because of their own risk-increasing behavior, but because of the predictable shape of others' errors.

This is unfair in a number of ways. When non-trivial risks of error are concentrated on a particular group rather than being randomly distributed, members of that group lose security. Even when violations are fully compensable these agents cannot have a reasonable expectation of being free from future violations, and so lose important values that moral rights are supposed to protect. But many violations are not fully compensable; it is better to avoid the violation than to suffer it and receive compensation. So, people who recognize themselves as belonging to a high-risk group sacrifice other goods in order to reduce their risk exposure, giving up activities and forms of dress or expression to which they had moral claim. Black men in the United States, for instance, recount wearing suits, avoiding ATMs after dark, whistling classical tunes if walking alone at night, or, more generally, trying to ‘avoid acting Black’, in hopes of avoiding being read as threat

²⁵I am intentionally leaving the precise characterization of subjective permissibility standards quite vague, to be consistent with a reasonably wide variety of proposals.

²⁶For data suggesting that, at least in American contexts, subjects appear prone to make precisely these errors, see for a start Correll, Wittenbrink, Park, Judd, and Goyle (2011); Eberhardt, Purdie, Goff, and Davies (2004); Richardson and Goff (2012) for threats, and Henningsen, Henningsen, and Valde (2006); Johnson, Stockdale, and Saal (1991); Osman (2003) for consent.

Objective Rights & Epistemic Risks 11

ening.²⁷ Members of these groups face three types of costs that they would be free from under conditions of perfect information: (a) they are disproportionately likely to suffer errors, which can be only imperfectly compensated; (b) they lack the security that rights are supposed to ensure, since even those who escape errors lack a reasonable expectation against suffering them; and (c) they suffer the additional losses associated with risk-reducing behaviors aimed at managing the first two costs. All of this is a predictable consequence of appealing to a subjective standard (invoking a threshold probability or expected moral value) to guide agents’ deliberation, given background facts about how social contexts shape agents’ rational credences.

This problem is worst for the evidence-relativist: Because their characterization of the objective demands of rights is fundamentally grounded in the decision-context of the duty-bound, the view implies that agents **lack rights** against sufficiently improbable errors. It lacks resources to acknowledge or correct for the fact that some people are disproportionately exposed to suffering such errors. This has ripple effects for compensation: Agents who suffer rights-violations have a claim to compensation even if the violation is excused. But if no right was violated, no wrong was done, and the victim lacks a claim to compensation even after the error is discovered. The Outcome-relativist faces a serious problem with distributive fairness too: though she can condemn actual errors as objective wrongings, she cannot affirm that individuals subjected merely to additional risks in (a), or to the costs (b) and (c), are thereby **wronged**, or that they have a moral claim to compensation for these burdens. So both views leave the disproportionately exposed with a less robust claim to the moral goods protected by core rights—in violation of the security and distributive fairness roles of rights.

To characterize their fair claims on others, rights under epistemic conditions like our actual circumstances must be both more inclusive and more limited than those implied by the ‘always’

and 'never' answers. In that we must recognize agents' rights against disproportionate risk exposure, they must be more inclusive; in that we cannot consider all errors to be equally wrongful, they must be less inclusive.

4.2 Failures of correlativity

Let's move on to evaluating how the duties implied by each standard align with the reasonable demands agents may make on each other, in virtue of their moral standing. Here the outcome-relative view falters. Because claim-rights generate duties, A's having a right that B ϕ entails that B has a corresponding duty, owed to A, to ϕ . But there are many cases where whether an agent has a duty, what that duty demands, and the relative priority of her duties seem to depend importantly on her epistemic position, and so **not** to correspond with rights as characterized by the outcome-relativist. This can be illustrated with a variant of Jackson (1991)'s **drug** example:²⁸

doctor — As a physician, Jill has duty of care to act in the best interests of her patient, John. John has a minor but not trivial skin complaint; there are 11 treatments available. Jill knows that treatment *a* will only partially cure him, but won't kill him. Of the remaining 10 treatments, Jill knows that exactly one would fully cure John, but she doesn't know which; and she knows that all the others will kill him. In fact *y* is the full cure. There are several things Jill could do:

- **Risky:** Reasoning that she has a duty to do what is in fact best for John, Jill tries, prescribing a non-*a* treatment, knowing (given the odds) it will almost certainly kill him. It does.

²⁷See generally Steele (2011). Similarly women may avoid expensive dates, dressing attractively when going out alone, etc., in order to avoid being read as consenting to sexual contact.

²⁸The original case Jackson (1991, p. 462-463.) was presented as an argument in support of a Subjectivist standard for right-action.

Objective Rights & Epistemic Risks 12

- **Lucky:** Jill reasons as before, and runs a $\geq 90\%$ risk of killing John. But luckily she prescribes *y*, and it cures him completely.
- **Safe:** Reasoning that it would wrong John to run a $\geq 90\%$ risk of killing him, Jill prescribes *a* knowing that it won't fully cure John, but also runs no risk of killing him. John is partially cured.

Jill knows that *a* is not the treatment a fully informed third party would advise her to prescribe, but insofar as John has a right that his physician act in his best interest, Jill violates this right if she performs either **risky** or **lucky**.²⁹ If her epistemic state makes it possible for her to safely effect a complete cure, she must do that; but if ignorant which of the treatments will achieve this, risking John's death rather than safely prescribing *a* would breach her duty. So even though **lucky** would in fact be best for John, in her actual epistemic position Jill can only fulfill her duty to act in John's best interest by performing **safe**. This implies (contra Outcome Relativism) that Jill's epistemic limitations are relevant to determining the content of her duty: the demands of the rights-based standard come apart from what a fully-informed advisor would recommend, and while what is **best** for her to do may be insensitive to her epistemic limitations, what she owes to John is not.³⁰

An outcome-relativist might—borrowing a defense against the parallel objection to Objectivism for right action³¹—try to accommodate this verdict by making the following appeal

to rights against risk:

Though **safe** does violate Jill's obligation to act in John's best interest, it satisfies a more stringent duty to avoid exposing him to unjustified risk of death (which **lucky** violates), and so **safe** is Jill's least bad option.

This move doesn't actually vindicate outcome-relativism. Though it denies that ignorance **extinguishes** duties the agent would have if fully informed, it concedes that ignorance is relevant to determining both the content and weight of her duties. There is nothing objectively chancy about the operation of the drugs in doctor; the risk is purely epistemic. Were Jill to know all the relevant facts, there would be no risks involved. So to even claim that Jill violates a stringent duty in **Lucky**, we must allow that facts about Jill's ignorance play a role in determining which duties she in fact has.

Outcome-relativism denies that epistemic states determine the contents of rights even in this limited way, forcing us to say that Jill blamelessly violates John's rights if she performs **safe**, and is blameworthy but does not wrong John if she performs **lucky**. There are two reasons to resist this. The first is that it is an awkward fit. Blameless wronging carves out a category in which the agent had good reason to think that what she was

²⁹It is inadvisable to try to characterize Jill's rights-based duties here by using phrases like "objectively, Jill ought..." or "subjectively, Jill ought...". The latter invokes the 'deliberative ought'—a subjective standard—but we want to know about Jill's objective duties. The former too easily invokes the 'evaluative ought', and it is of course true in the evaluative sense that Jill ought to prescribe *y*, but evaluative **oughts** track the comparative goodness of states of affairs, rather than the duties of agents. Neither reliably tracks agent's rights-based obligations. (For more discussion of the evaluative vs. agentive/prescriptive, and deliberative ought, see [Brennan & Southwood, 2007](#); [Schroeder, 2011](#)).

³⁰Thomson's preferred solution is to affirm that we objectively ought to avoid imposing high risks of harm on others, but deny that this is because any agent has a claim against being subjected to high risk of harm. Rather, the moral reasons to minimize expected harms gives us reason to regulate risk-imposing activities, and individuals objectively ought to comply with effective harm-minimization regulations (see 'Government Regulation of Behavior' and 'Imposing Risks' in [Thomson \(1986\)](#), and 'Harm' in [Thomson \(1990\)](#)).

³¹The reply (discussed by [Graham, 2010](#); [Mason, 2013](#)) the Objectivist can make to similar objections concerning overall moral obligation runs as follows: in Jill's situation, only prescribing *y* is objectively permitted by standards for right-action, but she is **more** morally obligated to avoid killing her patient than to avoid a merely partial cure, so she should (in the practical sense) prescribe *a*, violating one moral obligation in order to avoid violating a more stringent one.

Objective Rights & Epistemic Risks 13

doing was allowable, but was faultlessly ignorant of facts that made it a violation. But in **safe**, Jill is fully aware when she prescribes *a* that it is not what she owes to John according to the outcome-relative standard. The second is that it leaves us unable to distinguish the wrong done to John in **risky** from a case where the risks are absent, but the same harm results:

Mislabeled: Reasoning as before, Jill prescribes *a* believing it will partially cure John without risking killing him. Unknown to Jill, the labels were switched, and in administering the treatment labeled *a*, Jill actually gives John one of the other 10 treatments. It kills him.

John has a distinctive complaint against Jill in **risky** that he lacks in **mislabeled**. Jill's action causes John's death in both, but Jill gambled with his life only in the former. This difference between these cases cannot be reduced to a difference in Jill's degree of blameworthiness; what distinguishes the cases is a fact about the respect John is owed in virtue

of his moral status.³²

So the first way that outcome-relativism intuitively mischaracterizes duties is that it cannot recognize the duties an agent incurs precisely **because** she is ignorant. A second is that it severs the link between rights and accountability—decoupling having rights from having standing to insist on fulfillment or issue a reproach for a breach—because it yields duties which are unreasonable in at least two respects:

- (i) it requires the duty-bound to shoulder unjustifiably high epistemic risks. Suppose that rather than killing John, the bad treatments in doctor would somehow kill Jill. A fully-informed advisor would still tell Jill to choose *y*, knowing it will in fact both fully cure John and not kill Jill. However, since **given her actual epistemic position** doing so requires Jill to run a $\geq 90\%$ (epistemic) risk of dying in the attempt, this is not something that John can reasonably demand that she do.
- (ii) it condemns behavior that is irreproachable by its own lights. As [Ferzan \(2005\)](#) and [Quong \(2015\)](#) point out, outcome-relativism implies that when a justified risk-imposing activity like conscientious driving results in the harm risked, the agent has failed in a duty, violating the right of individual who suffers the harm—but there is no reproach the sufferer could address to the agent. The sufferer cannot demand that the driver be more careful, or abstain from driving. After all, the activity is justified: the benefits of the practice outweigh the risks, and the conscientious driver has already done all she can to minimize the risk. The duty breached could only be a duty to **not drive when doing so will actually, unpredictably, result in harm**; and the driver is in principle incapable of having the information necessary to ensure that her driving does not breach this duty.

The final major problem with the duties implied by outcome-relativism is that they are counterfactual: the duties an agent **would** have, **if she knew all the relevant facts**. But for a whole range of features, it is inappropriate to abstract away from an agent's actual circumstances when determining what her actual duties are. The fact that **if** Bea were impervious to flame she would have a duty to rescue Abe from a burning building does not tell us anything about whether Bea **as she actually is** has such a duty. That depends on what it will cost Bea to rescue Abe, given her actual physical condition. We do not in general determine an agent's duties by imagining what duties she **would** have, were she not subject to the various circumstantial constraints to which she is in fact non-culpably subject. Why then should we think that the duties Bea would have **were she fully-informed** are a good guide to her actual duties when non-culpably ignorant? In abstracting away

³²Compare [Lazar \(2017\)](#).

Objective Rights & Epistemic Risks 14

from agents' actual epistemic limitations, outcome-relativism treats them as radically different from agents' other limitations, without offering justification for doing so.³³

In sum, the outcome-relative characterization of objective rights does not correlate well with the duties that are supposed to arise from those rights. It excludes the duties an agent incurs

precisely **because** she is ignorant, severs the connection between having a right and being in a position to demand performance or reproach a breach, and singles out epistemic limitations as the only agential limitations that have no effect on the content, stringency, and priority of our duties. An evidence-relative standard does **a little** better: it gives the right verdicts in cases like doctor, it does not posit unreasonable duties, and it does take epistemic risks to be relevant to rights. But it falters elsewhere by denying that duties correlative agents' rights (to property, or against harm) can persist despite misleading evidence, yielding the wrong verdicts in the coat case and the opening case.

4.3 The deliberative role

We've come finally to the role of rights in guiding deliberation. Some ways of articulating the objection that outcome-relativism cannot guide deliberation strongly resemble the parallel objection that Objectivist standards of right-action are not 'useable'. There is a strong temptation to think that the replies offered there can be leveraged to similar effect in the present context. But here we must go very carefully. It is uncontroversial that a standard of right action is evaluative: it enables us to keep the moral score, ruling on whether (retrospectively or from the point of view of the universe) an action was right. Many argue that it must also play a role in guiding action, for which it is necessary that the standard be "in some appropriate sense be present to the agent's mind."³⁴ As Smith (2012) reconstructs this objection, a standard for right-action is inadequately action-guiding if agents are unable to **use** the principle to decide what to do. Objectivists respond that the standard need not do this work directly, but can instead be supplemented by a subjective decision rule to guide agents who aim to do the right thing but suffer from uncertainty.³⁵ This renders the standard useable in the following sense: the agent can choose her actions for the reasons that the standard identifies as right making by following the prescriptions of the decision rule. When (because of her ignorance) this leads her to do something forbidden by the objective standard, she has done wrong but is excused.³⁶

Rights work differently. They are not just esoteric facts about the moral scoreboard; they are the operative rules of the game **For the players**. Moral rights issue relational rules, structuring the real-time interactions of agents in their actual circumstances. They give rise to duties not just to secure certain results, but to deliberate in certain ways. A standard that is insensitive to agent ignorance cannot characterize rights in a way that will fill these roles. First, they must be the sorts of things that are generally available to be taken as reasons in agents' **ex ante** deliberation, in their actual circumstances. That Abe has a right against being harmed is a reason which can guide Bea's deliberation only if she is in a position to recognize which of her available actions it affects.³⁷ When (like Jill in doctor) ignorance blocks her from knowing whether a contemplated

³³Arguments along the lines of the ones presented in this section also appear in Zimmerman (2008), Quong (2015), and Quong (2020).³⁴Jackson (1991, p. 467)

³⁵Driver (2013), Graham (2010), and Feldman (2012) offer arguments to this effect.

³⁶See Smith (2018, **Making Morality Work**) for an extended development and defense of a Hybrid

approach of this kind.³⁷Though giving a full characterization of the relevant sense of 'available' demands a paper-length treatment of its own, I'll mention a

action would constitute harming, she cannot guide her deliberation by reference to the fact of Abe's right. It is not sufficient that it is **physically** possible for her avoid harming. Second, whether Bea satisfies her duty to give appropriate deliberative weight to Abe's interests is a function not of the actual results of her action, but of the epistemic profile of her choice.³⁸ It is a matter of what alternatives she considered, what evidence she sought, and what risks she took herself to be running—even if, like Jill's, those risks were purely epistemic. It follows that facts about her epistemic circumstances are relevant to whether she has fulfilled her duty, and so must be relevant to what the right (which grounded her duty) demanded of her.

A supplemental subjective standard can provide practical advice, but unlike right-action, the deliberative problem for standards of rights isn't just that agents don't know **what to do**. It's that others' rights must figure directly in the duty-bound's deliberation in particular ways, and that some rights have essentially deliberative contents. A subjective standard cannot deliver the right reasons for action; knowing that an action is subjectively permissible only assures Bea that it is **either** not a violation or is excused. And that an action would be excused is no reason to choose to perform it, and taking this as a guiding reason is a far cry from taking Abe's rights as constraining one's deliberation about what to do.³⁹

This is not to say that a supplemental standard would have no use— it can help guide agential deliberation in hard cases, though there are many reasons why the objective standard must still not diverge too far from what agents could usually track given their actual epistemic limitations. Whether a standard secures moral goods for rights-holders depends on how other agents guide their behavior. The accountability of duties and the deliberative functions of rights will not tolerate too great a divergence between what the rule demands of agents and what agents are usually in a position to at least guess. And since the moral value of an action **vis à-vis** rights is partly constituted by the conditions under which it is made—including the revealed attitude toward various risks—epistemic risks can be relevant to determinations of what we duties we have, not merely to whether we are blameworthy.

But though simply supplementing an outcome-relative standard with subjective guidance is inadequate to fulfill these functions, they do not quite require that duty-bound agents always be in a position to know what their duties require of them **no matter their epistemic circumstances**. deliberative guidance does not push us all the way to an evidence-relative standard.

5 Wrapping up: Nuanced Fact-Relative Rights

We could insist on retaining outcome-relativism because it is tidy. Or because we take it to be the only alternative to an evidence-relative view. But I hope to have shown that if we thoroughly exclude agent

couple of options before moving on. One way we could go is to say that ϕ ing is available to A if it is something she can **intentionally do**, where one can intentionally perform an action under a description only if one knew or could reasonably guess at the time of action that the act satisfies the description. Another alternative (developed in [Cariani, Kaufmann, & Kaufmann, 2013](#)) is to say that an action is available only if it is 'choosable', where ϕ is choosable iff there's some specific action β such that it's epistemically necessary for A that she can perform β , and that performing β entails achieving ϕ . Another approach (defended in [Hedden, 2012](#)) counts ϕ ing as an option for A at t only if ϕ ing is something A can decide to do and can bring about at t .

³⁸This holds whether we couch it as the claim that agents have rights against purely epistemic risks (see Finkelstein, 2003; Oberdiek, 2008), or to robust goods like security and respect (see Kumar, 2003; Lazar, 2017; Pettit, 2015; Quong, 2012; Raz, 1986), or as the claim that a certain deliberative role is constitutive of all moral rights (see Darwall, 2006; T. Scanlon, 1998).

³⁹See also Rodin (2003, *War and Self-Defense*), pp. 30-31.

Objective Rights & Epistemic Risks 16

ignorance from affecting what is allowable, there's almost nothing left of the roles and interpersonal structure distinctive to moral rights. This leaves us with two options: give up outcome-relativism, or give up the roles it can't fill. Thomson (1990) sacrificed the roles, denying that rights bind only moral agents or play an important deliberative role, rejecting the existence of positive rights and rights against risks, and undermining reproach for breaches of duty. This effectively reduces talk of rights from constraints on how agents should (in the prescriptive sense) treat each other to a mere subcategory of the evaluative ought—an exceptionally high price for a contemporary rights-theorist to pay. It seems to me that the rich theoretical roles of rights are worth preserving, and we should seek a better characterization. The inadequacies we've discussed help to sketch the profile of a better way.

To get the right verdict in the police case, objective rights cannot be a pure function of the duty-bound's epistemic position. Rights that come and go with changes in others' evidence fail security, compensation, distributive fairness, and correlativity. So we must take the demands of objective rights to be characterized by facts, rather than by the agent's evidence. But to get the right verdict in the doctor case, and more generally to satisfy the deliberative and correlativity constraints on rights, we must accept that agent-ignorance has some bearing on what objective duties she has, not merely to whether she is blameworthy. The most conservative way to do this accepts the basic outcome-relative characterization of rights—we have all the duties we would have, were we to know all the relevant facts—but then extends it, allowing that epistemic risks can make infringements **morally worse**, and ground additional duties.⁴⁰ This accommodates the intuition that John has an extra complaint in **risky**. However, it still mischaracterizes objective duties: it still requires agents like Jill to act in accordance with facts they cannot track, and to accept unreasonable epistemic risks, and still implies the existence of some duties which agents are **in principle** irreproachable for violating (like the conscientious driver). It also does little to address the failures of security.

So we should go further, allowing that facts about the agent's epistemic position sometimes change whether an error counts as a violation at all.⁴¹ This means holding that some facts about agents' epistemic limitations are relevant to determining whether a duty exists, while others are relevant only to our evaluation of an agent's culpability for breach of that duty. This leaves us with the interesting and important project of determining which is which, and there are many options. An obvious one is to define a minimal probabilistic threshold, such that errors with an evidence-relative probability below the threshold are not violations.⁴² But neither treating ignorance as wholly irrelevant nor uniformly discounting by evidential probability can satisfy distributive fairness under our actual social conditions. To thread the needle, the rights-based standard must discriminate between various sources of error in a way that restores a reasonable expectation against suffering violations to members of groups that are disproportionately exposed to risk of error.

It is worth emphasizing that the resultant duties are not necessarily cognitively demanding, nor are their contours esoteric. Members of disproportionately exposed groups have a moral claim that you and others manifest equal regard for their rights and security as for others' when deliberating. Under ideal conditions,

⁴⁰This is one way of reading the main argument in Lazar (2017).

⁴¹This is not a lonely path—many theorists have suggested that what rights others have on our behavior must be constrained in some way by what evidence is or could reasonably be available to us. A haphazard collection of advocates this include Kumar (2003), Oberdiek (2017), Quong (2015), Ripstein (1996), T. Scanlon (2008), and Tadros (2011, chpt 10).

⁴²This is suggested by McMahan (2009)'s version of the outcome-relative view, which he restricts to 'foreseeable' harms.

Objective Rights & Epistemic Risks 17

you could do this by simply being guided by your actual rational credences in all cases. But under non-ideal conditions—when your evidence is partly shaped by upstream social injustices—respecting their right to be equally secure in their expectation against suffering violations will require you to refrain from contributing to systems or patterns of action that deprive oppressed groups of such security. It is not particularly cognitively demanding to know about this duty, nor would you need substantial sociological knowledge to know what it demands from you. From the armchair, you can know that agents have a moral claim against being deprived of security on the basis of their membership in oppressed social groups. With the addition of minimal empirical information you can know that members of groups that are over-represented in crime statistics will be mistakenly read as threats disproportionately often. So, a good-faith effort to avoid contributing to a pattern of action that violates the group-members' claim to security will require that you not invoke these statistics to help justify running a higher risk of mistakenly harming them. This is not a particularly complicated duty: it's just a diachronic version of the obligation to guide one's deliberation in ways that avoid imposing excessive epistemic risks on others.⁴³

We should at least be open to the possibility that the ways in which agent ignorance is relevant to the demands of objective rights will not, in the end, be captured by any tidy picture or decision-theoretic function, but will instead partly depend on what values the right protects for an agent. Plausibly, some rights—rights requiring others to act in one's best interest, perhaps—will be wholly evidence-relative, while others—ones which serve to mark the agent as a moral person with claims equally stringent as others'—will track the behavior of the **rights-holder**.⁴⁴ These are just indicative starting points; I am not here offering a fully worked out alternative. I have only sought to demonstrate that there is space for one; an agent's epistemic situation is sometimes relevant to what she must do to respect others' rights, and hence affects objective allowability, and we can acknowledge this while retaining a commitment to taking objective rights to be fact-relative.

Though these were humble ambitions, they suffice to disrupt the condensed argument; we can affirm that the resident was wronged while denying that the same would hold of a bluffing or ineffectual gunman. The methodological upshot is also significant: if agents' ignorance sometimes affects what is objectively allowable, then we cannot theorize about the demands of agents' actual moral rights by looking only at what outcomes they have an interest in, or by

imagining what rights would require from fully-informed agents. Nor can we infer from the fact that a consideration is epistemic that it does not affect whether an action is a rights violation, or can safely be relegated exclusively to the domain of subjective permissibility. Once we give up the assumption that agent-ignorance affects only blame, the terrain is not so tidy as we might have thought.

In short, there is a lot of work left to do.

⁴³My thanks to Kieran Setiya for a helpful conversation drawing out these points.

⁴⁴One proposal along these lines, floated by Quong (2015) holds that if the subjective permissibility of A's ϕ ing depends on false moral beliefs (including whether B has altered his rights), A's ignorance only excuses, and otherwise it alters her objective duties. Bolinger (2021) discusses another alternative, on which when misleading evidence arises from A's responsible agency, it undermines their complaint against an error, and so can release B from an objective duty. This is contrasted with other forms of evidence, which is glossed as only affecting B's culpability. I offer a more fully worked-out account in *Rewriting Rights* (ms), arguing that rather than characterizing rights either by reference to outcomes for A or by reference to prospects from B's evidential perspective, we should frame them as obligations to **pursue particular cooperative strategies**. And of course there are plenty of other options.

Objective Rights & Epistemic Risks 18

References

Bolinger, R. J. (2021). The moral grounds of reasonably mistaken self-defense.

***Philosophy and Phenomenological Research*, 103**(1), 140-156.

Brennan, G., & Southwood, N. (2007). Feasibility in action and attitude. In T. Rønnow-Rasmussen, B. Petersson, J. Josefsson, & D. Egonsson (Eds.), ***Hommage à Wlodek. Philosophical Papers Dedicated to Wlodek Rabinowicz***. (Published as web resource only).

Cariani, F., Kaufmann, M., & Kaufmann, S. (2013). Deliberative modality under epistemic uncertainty. ***Linguistics and Philosophy*, 36**(3), 225-259.

Cornell, N. (2015). Wrongs, rights, and third parties. ***Philosophy and Public Affairs*, 43**(2), 109-143. Correll, J., Wittenbrink, B., Park, B., Judd, C., & Goyle, A.

(2011). Dangerous enough: Modeling racial bias with contextual threat cues. ***Journal of Experimental Social Psychology*, 47**, 184-189.

(doi:10.1016/j.jesp.2010.08.017)

Darwall, S. (2006). ***The second person standpoint: Morality, respect, and accountability***. Harvard University Press.

Driver, J. (2013). What the objective standard is good for. In M. Timmons (Ed.), ***Oxford studies in normative ethics*** (Vol. 2). New York: Oxford University Press.

Dworkin, R. (1978). ***Taking rights seriously***. Harvard University Press.

Eberhardt, J. L., Purdie, V. J., Goff, P. A., & Davies, P. (2004). Seeing black: Race, crime, and visual processing. ***Journal of Personality and Social Psychology*, 87**(6), 876-893.

Feinberg, J. (1966). Rights, duties, and claims. ***American Philosophical***

- Quarterly**, 3(2), 137-144. Feldman, F. (2012). True and useful: On the structure of a two level normative theory. **Utilitas**, 24(2), 151-71. Ferzan, K. K. (2005). Justifying self-defense. **Law and Philosophy**, 24(6), 711-749. (doi: 10.1007/s10982-005-0833-z)
- Finkelstein, C. (2003). Is risk a harm? **University of Pennsylvania Law Review**, 151(3), 963-1001. Foot, P. (1967). The problem of abortion and the doctrine of double effect. **Oxford Review**, 5, 5-15. Frowe, H. (2010). **A practical account of self-defence** (Vol. 29) (No. 3). (doi: 10.1007/s10981-009-9062-1)
- Frowe, H. (2015). Claims, duties, and lesser-evil justifications. **Proceedings of the Aristotelian Society**, 89(1), 267-285.
- Graham, P. (2010). In defense of objectivism about moral obligation. **Ethics**, 121, 88-115.
- Hedden, B. (2012). Options and the subjective ought. **Philosophical Studies**, 158(2), 343-360. Henningsen, D., Henningsen, M. L. M., & Valde, K. S. (2006). Gender differences in perceptions of women's sexual interest during cross-sex interactions: An application and extension of cognitive valence theory. **Sex Roles**, 54, 821-829.
- Hohfeld, W. N. (1919). **Fundamental legal conceptions as applied in judicial reasoning** (W. W. Cook, Ed.). Yale University Press.
- Hurd, H. (1999). Justification and excuse, wrongdoing and culpability. **Notre Dame Law Review**, 74(5), 1551-1574.
- Jackson, F. (1991). Decision-theoretic consequentialism and the nearest and dearest objection. **Ethics**, 101(3), 461-482.
- Jackson, F., & Smith, M. (2006). Absolutist moral theories and uncertainty. **Journal of Philosophy**, 103(6), 267-283.
- Johnson, C., Stockdale, M., & Saal, F. (1991). Persistence of men's misperceptions of friendly cues across a variety of interpersonal encounters. **Psychology of Women Quarterly**, 15, 463-475.
- Kramer, M. H. (2001). Getting rights right. In M. H. Kramer (Ed.), **Rights, wrongs and responsibilities** (p. 28-95). Pallgrave Macmillan.
- Kumar, R. (2003). Who can be wronged? **Philosophy and Public Affairs**, 31(1), 99-118. Lazar, S. (2017). Risky killing: How risks worsen violations of objective rights. **Journal of Moral Philosophy**, 1-26.
- Mason, E. (2013). Objectivism and prospectivism about rightness. **Journal of Ethics and Social Philosophy**, 7(2), 1-21.
- McMahan, J. (2005). The basis of moral liability to defensive killing. **Philosophical Issues**, 15(1), 386-405. (doi: 10.1111/j.1533-6077.2005.00073.x)
- McMahan, J. (2009). **Killing in war**. Oxford University Press.
- Moore, G. E. (1903). **Principia ethica**. Cambridge: Cambridge University Press.
- Nozick, R. (1974). **Anarchy, state, and utopia**. Basic Books.
- Oberdiek, J. (2004). Lost in moral space: On the infringing/violating distinction and its place in

- the theory of rights. **Law and Philosophy**, **23**, 325-346.
- Oberdiek, J. (2008). Towards a right against risking. **Law and Philosophy**, **28**(4), 367–392. doi: 10.1007/ s10982-008-9039-5
- Oberdiek, J. (2017). **Imposing risk: A normative framework**. Oxford University Press.
- Osman, S. L. (2003). Predicting men's rape perceptions based on the belief that 'no' really means 'yes'. **Journal of Applied Psychology**, **33**(4), 683-692.
- Parfit, D. (2011). **On what matters** (Vol. I). Oxford University Press.
- Pettit, P. (2015). **Robust demands of the good**. Oxford University Press.
- Quinn, W. (1989). Actions, intentions, and consequences: The doctrine of doing and allowing. **The Philosophical Review**, **98**(3), 287-312.
- Quong, J. (2012). Liability to defensive harm. **Philosophy and Public Affairs**, **40**(1), 45-77. (doi: 10.1111/j.1088-4963.2012.01217.x)
- Quong, J. (2015). Rights against harm. **Proceedings of the Aristotelian Society**(89), 249-266. (doi: 10.1111/j.1467-8349.2015.00252.x)
- Quong, J. (2020). **The morality of defensive force**. Oxford: Oxford University Press.
- Raz, J. (1986). **The morality of freedom**. Clarendon Press, Oxford.
- Richardson, L. S., & Goff, P. A. (2012). Self-defense and the suspicion heuristic. **Iowa Law Review**, **98**, 293-336.
- Ripstein, A. (1996). Self defense and equal protection. **University of Pittsburgh Law Review**, **57**, 685-724.
- Ripstein, A. (1999). **Equality, responsibility, and the law**. Cambridge University Press.
- Rodin, D. (2003). **War and self-defense**. Clarendon Press.
- Ross, D. (1930 [2002]). **The right and the good** (P. Stratton-Lake, Ed.). Clarendon Press, Oxford.
- Ross, D. (1939). **Foundations of ethics**. Oxford: Oxford University Press.
- Scanlon, T. (1998). **What we owe to each other**. The Belknap Press of Harvard University.
- Scanlon, T. (2008). **Moral dimensions**. Cambridge University Press.
- Scanlon, T. M. (2003). **The difficulty of tolerance: Essays in political philosophy**. Cambridge: Cambridge University Press.
- Schroeder, M. (2011). Oughts, agents, and actions. **The Philosophical Review**, **120**, 1-41.
- Objective Rights & Epistemic Risks** 20
- Smith, H. M. (2012). Using moral principles to guide decisions. **Philosophical Issues**, **22**, 369-386.
- Smith, H. M. (2018). **Making morality work**. Oxford: Oxford University Press.
- Steele, C. M. (2011). **Whistling vivavadi: How stereotypes affect us and what we can do**. W.W. Norton & Company.
- Tadros, V. (2011). **The ends of harm**. Oxford University Press.
- Tadros, V. (2016). Causation, culpability, and liability. In C. Coons & M. Weber (Eds.), **The**

ethics of self defense (p. 110-130). Oxford University Press.
Thomson, J. J. (1986). **Rights, restitution, and risk** (W. Parent, Ed.). Harvard University Press. Thomson, J. J. (1990). **The realm of rights**. Harvard University Press.
van der Vossen, B. (2016). Uncertain rights against defense. **Social Philosophy and Policy**, 32(2), 129-145. Zimmerman, M. (2006). Is moral obligation objective or subjective? **Utilitas**, 18(4), 329-61. Zimmerman, M. (2008). **Living with uncertainty**. Cambridge University Press.