Copyright

by

Samuel Fox Krauss

2019

**The Dissertation Committee for Samuel Fox Krauss Certifies
that this is the approved version of the following Dissertation:**

Three Essays on Applied Epistemology

Committee:

Sinan Dogramaci, Supervisor

Joshua Dever

Ernest David Sosa

Julia Staffel

Three Essays on Applied Epistemology

by

Samuel Fox Krauss

Dissertation

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

The University of Texas at Austin

May 2019

Dedication

I dedicate this dissertation to my parents.

Rabbi Jose ben Kisma said, "I was once walking by the way when a man met me and greeted me, and I returned his salutation. He said to me, 'Rabbi, from what place are thou?' I said to him, 'I come from a great city of Sages and Scribes.' He said to me, 'If thou art willing to dwell with us in our place, I will give thee a thousand thousand golden denars and precious stones and pearls.' I replied to him, 'If thou wert to give me all the silver and gold and precious stones and pearls in the world, I would not dwell elsewhere but in a place of Torah.'"

And thus it is written in the *Book of Psalms* by the hands of David, King of Israel, "The Law of Thy mouth is better unto me than thousands of gold and silver, and not only so, but in the hour of a man's departure, neither silver nor gold nor precious jewels nor pearls accompany him, but only Torah study and good works, as it is said, 'When thou walkest it shall lead thee, when thou liest down it shall watch over thee, and when thou awakest it shall talk with thee.'"

—Avot 6:9

Acknowledgements

Abstract

Three Essays on Applied Epistemology

Samuel Fox Krauss, Ph.D.

The University of Texas at Austin, 2019

Supervisor: Sinan Dogramaci

This dissertation contains three essays, each of which discusses a distinct way in which the particular status our beliefs have should affect the way we treat others.

In the first essay, I begin with an account of *epistemic damage*—a proposal about how to measure epistemic harm. Second, I give an account of *expected epistemic damage*, which allows us to draw a principled line from epistemic harm to moral blameworthiness. Third, I use the notion of expected epistemic damage to solve a dilemma I pose for the dominant account of lying. I critique, and offer a replacement for, a widely-accepted necessary condition on lying—that the speaker believes the negation of what they assert.

In the second essay, I argue that when people behave in a way that we believe is morally impermissible but toward which they are morally indifferent, we ought to pay them to forgo that behavior. People have legal entitlements to act in some ways that others regard as morally impermissible. But people exercise these entitlements, nevertheless. When they do, others have a defeasible reason to stop them. The circumstances will dictate whether they should, and, if so, the best method: one might convince them that what they are doing

is wrong; one might explain that people will dislike them if they persist; one might ask them nicely, or threaten them. Or, one could pay them.

In the third essay, I address arguments in both the philosophical and legal literature according to which statistical evidence cannot alone be sufficient evidence for a judgment in a civil trial or a conviction in a criminal trial. I argue that this dominant view is mistaken. Broadly, the argument relies on the presumption that any probative evidence ought to be given its due. I argue that the very many arguments presented against the sufficiency of statistical evidence are not strong enough to overcome this presumption.

Table of Contents

1 INTRODUCTION

This dissertation contains three essays, each of which discusses a distinct way in which the particular status our beliefs have should affect the way we treat others.

In the first essay, I begin with an account of *epistemic damage*—a proposal about how to measure epistemic harm. Second, I give an account of *expected epistemic damage*, which allows us to draw a principled line from epistemic harm to moral blameworthiness. Third, I use the notion of expected epistemic damage to solve a dilemma I pose for the dominant account of lying. I critique, and offer a replacement for, a widely-accepted necessary condition on lying—that the speaker believes the negation of what they assert. The hope is that this will bring some clarity and precision to a frequent kind of interaction and make it easier to assess damage and assign blame. The contribution of the paper is to argue that the degree of confidence we have in what we assert, in part, determines the extent to which we are blameworthy for certain assertions.

In the second essay, I argue that when people behave in a way that we believe is morally impermissible but toward which they are morally indifferent, we ought to pay them to forgo that behavior. People have legal entitlements to act in some ways that others regard as morally impermissible. But people exercise these entitlements, nevertheless. When they do, others have a defeasible reason to stop them. The circumstances will dictate whether they should, and, if so, the best method: one might convince them that what they are doing is wrong; one might explain that people will dislike them if they persist; one might ask them nicely, or threaten them. Or, one could pay them. When our beliefs match up with others in a particular way—where we believe they are behaving impermissibly and they

believe they are engaging in permissible but not obligatory behavior, the particular mis-match of the attitudes create an opportunity for moral improvement.

In the third essay, I address arguments in both the philosophical and legal literature according to which statistical evidence cannot alone be sufficient evidence for a judgment in a civil trial or a conviction in a criminal trial. I argue that this dominant view is mistaken. Broadly, the argument relies on the presumption that any probative evidence ought to be given its due. I argue that the very many arguments presented against the sufficiency of statistical evidence are not strong enough to overcome this presumption.

2 EPISTEMIC DAMAGE

If someone punches you in the nose, we have a way to assess the harm to the body. And we can determine legal liability and moral blameworthiness. When you are misled, by contrast, you suffer *epistemic* harm, in addition to other harms you might suffer. But, we don't quite have a way to assess epistemic harm or a principled way to relate epistemic harm to moral blameworthiness. This paper proposes to do both. First, I give an account of *epistemic damage*, which is a proposal about how to measure epistemic harm. Second, I give an account of *expected epistemic damage*, which allows us to draw a principled line from epistemic harm to moral blameworthiness. Third, I use the notion of expected epistemic damage to solve a dilemma I pose for the dominant account of lying. The hope is that this will bring some clarity and precision to a frequent kind of interaction and make it easier to assess damage and assign blame. I'll note from the outset that much of this project is exploratory.

Imagine that Albert, but not Bob, knows that the forecast says it will rain. Before he leaves the house, Bob asks Albert whether it will rain. Albert lies to Bob and tells him that it won't. Bob thereby comes to believe that it won't rain, leaves the house for the day without an umbrella, and later gets soaking wet and catches pneumonia. Clearly, Albert has done something wrong, and is at least partially responsible for Bob's getting wet, if not also for his getting sick. But he also inflicted another kind of harm on Bob—*epistemic harm*: Albert caused Bob to come to have a false belief that Albert knew was false. How should we assess this harm and how should we determine the extent to which Albert is blameworthy? To answer these questions, we have to ask, first, what is the object of the

3

damage, and, second, in virtue of what is that thing damaged? Third, we'll have to figure

out the relationship between epistemic harm and moral blameworthiness. In what follows,

I'll use what I take to be a simple, paradigm case of infliction of epistemic damage: lying.

In this case, the speaker lies to the *dupe*[1], who thereby comes to (falsely) believe the

proposition that the speaker asserted.

2.1 Epistemic Damage

2.1.1 INTRODUCTION

The first step is to determine what, or who, suffers epistemic damage. The object

(victim) of the damage, I want to say, is the person. That is, while parts of a person might

be said to have been damaged, in the relevant sense the kind of things one can be

blameworthy for damaging, at least in this context, are *people*. This is less than obvious. If

someone is punched in the nose, we might give several different answers to the question

"What (or who) is the object of the damage?" We might say that the nose was damaged, or

the nasal septum. But in determining liability, it seems to me the relevant fact is that the

person was harmed. And while we can determine the extent to which they were harmed by

looking at their nose, the interesting thing is that *they* were harmed, not their nose, and not

their septum, even though we might speak, informally, that way. In the same way, I'll

assume from here that the object (or victim) of the epistemic damage is the person who's

been harmed, and not their *beliefs* or *mind* or anything else. I'm not saying that this is *the*

---

[1] As best I could determine, the first use of *dupe*, in the philosophical literature, is due to Bok, *Lying: Moral Choice in Public and Private Life*.

correct view of the object of damage in every case, but I'm taking it for granted because it makes the most sense for my purposes.

The second step, which is really the task of this section, is to determine the form the damage takes. Or, to put it another way, to determine *in virtue of what* someone can be truthfully said to have been epistemically damaged. To determine this we have to first determine what is epistemically *important* (or *valuable*). Here are some contenders: full beliefs; accuracy; knowledge; justification; the total number of beliefs; important beliefs; epistemic abilities (or rational capacity).

My sense is that practical considerations will at least partially determine which of the above is the most salient, and this will vary by context. Sticking with the nose example, we might ask what it is for a person to suffer damage to their nose. Well, one way is to make it such that the person can no longer smell. To assess the extent of this kind of damage, we can administer smelling tests: the greater the loss of smell, the greater the damage. However, in some situations we won't care about olfactory capabilities, and in those cases we would want to talk about the structure of the nose: we'd want to know which bones were broken, and how badly. It doesn't seem as though there is a fact of the matter about which *kind* of nasal damage is the most important, but rather one or the other might be more salient. It could be that there's one best, all-things-considered kind of damage. Here I'll articulate what I think is the most compelling, all-purpose account of epistemic damage, although I do not mean to suggest that this is the only, or best account. Then, I'll briefly describe how we might use other contenders in the account of epistemic damage.

2.1.2 INCREASE IN INACCURACY AS EPISTEMIC DAMAGE[2]

Philosophers who theorize in terms of full beliefs may think about the blameworthiness of liars in the following way. A rational agent uses her beliefs, in combination with her desires, to act in order to satisfy her desires.[3] When a person has true beliefs, she acts in ways that tend to satisfy her desires; when a person has false beliefs, she acts in ways that tend to frustrate her desires. Thus, we care about having true beliefs. (This is consistent with acknowledging that it may not be the case that it's *always* better for us, practically, to have true beliefs.[4]) One might instead endorse a non-instrumentalist view about the value of true beliefs. I won't discuss this view further, but it is consistent with what I say below.

A degree-theoretic account of belief, however, allows for greater precision and a finer explanation. We can say, instead, that a person has a *credence* in a proposition, where a credence is a subjective probability one assigns to the truth of a proposition and ranges from 0 (the minimal degree of certainty) to 1 (the maximal degree of certainty.) While beliefs are either true or false, credences are more or less *accurate*. Accuracy is measured by a scoring rule, which takes as its inputs the agent's credence and the truth of the proposition. We might think of the accuracy of a credence as a measure of "its 'closeness

---

[2] This section, as well as Sections 2.2.1, 2.3.1, 2.3.2, and 2.3.3, are nearly identical to portions of Krauss, Sam Fox, "Lying, risk and accuracy." *Analysis* 77(4) (2017): 726–34, of which I am the sole author.

[3] *See*, e.g., Ramsey, *Truth and Probability*, 68–74; Stalnaker, *Inquiry*, 18; Dennett, *The Intentional Stance*, 20–33; 43–57.

[4] *See*, e.g., Kelly, "Epistemic Rationality as Instrumental Rationality: A Critique."

to the truth.'"[5] For example, certainty of a true proposition is maximally accurate and certainty of a false proposition is maximally inaccurate. Being unsure of a true proposition is more inaccurate than if one were sure, but less inaccurate than if one were sure of its negation. I'll note that there is significant controversy about the appropriate desiderata for epistemic scoring rules.[6] For purposes of this paper, though, we can abstract away from most issues that cause controversy. We need only assume what is widely agreed upon: that accuracy scores of credences increase as the credence approaches the truth value of the proposition.

Because we need accurate credences to act in ways that tend to satisfy our desires, we care about having accurate credences. Call the increase in inaccuracy the victim suffers in any discrete interaction *epistemic damage*.

Consider an instance of lying from the perspective of an omniscient third party: the speaker has some small credence in a false proposition that she nevertheless asserts to the dupe. If the dupe trusts the speaker, the dupe will count the speaker's assertion as evidence for the proposition asserted, and the dupe's credence in that proposition will increase. Because the dupe has become more confident in a false proposition, she is epistemically worse off—she has suffered epistemic damage. The greater the inaccuracy of our credences, the less we will tend to get what we want. So, epistemic damage correlates with

---

[5] Schoenfield, "An Accuracy Based Approach to Higher Order Evidence," 694.

[6] See the literature on epistemic scoring rules, e.g., Joyce, "A Nonpragmatic Vindication of Probabilism"; Schoenfield, "Bridging Rationality and Accuracy"; Konek and Levinstein, "The Foundations of Epistemic Decision Theory."

real-world, or, *practical* harm. (This is consistent with acknowledging both that epistemic damage does not always cause practical harm and that lying does not always lead to epistemic damage.)

2.1.3 ALTERNATIVE ANSWERS TO THE QUESTION "IN VIRTUE OF WHAT IS THE PERSON DAMAGED?"

In the previous section, I gave an account of epistemic damage using accuracy because I think it is the most useful. Although I think an accuracy-first picture makes sense, I don't want epistemic damage to be limited to it. So, although I'll discuss epistemic damage in terms of accuracy, I'll below elaborate a few other ways in which we could think of it.

We might want to talk about epistemic damage as concerning full beliefs. Thus, we could say something like the following: one is epistemically damaged if they are made to have a false belief or to suspend a true belief. Here, if we are talking about full beliefs, there are three possibilities: the victim was agnostic with respect to the proposition and was convinced of the false proposition's truth; the victim believed the true proposition and became agnostic; or the dupe was agnostic and came to believe its negation. It seems to me that the dupe suffered the greatest epistemic damage in the last case. With full beliefs, however, it's not entirely clear how *confidence* gets into the picture. What if the audience believes a proposition, the speaker tells them that it's false, and while they retain the belief, they become less confident in it? On a full-belief picture, it's not clear what changes, or what effect on practical deliberation there will be. In any case, someone who rejects the

degree of belief model might be inclined to accept something like this picture, which, although I don't favor it, is not inconsistent with the account I want to give.

Epistemic damage might concern knowledge. Imagine that someone knows a proposition, but then, through someone else's action, loses that knowledge. This could be in virtue of a loss of any of the constitutive parts of knowledge. Say, for example, someone knows that the restaurant closes at 10pm; an ill-intentioned Albert could epistemically damage Bob by making him cease to know that about the restaurant. Albert could, for example, lie to Bob by telling him that it closes at 9pm, instead. If we imagine that this changes the dupe's belief, so that he no longer believes that it closes at 10pm, then he no longer knows it. Or, even if dupe doesn't change his belief, it could be that the speaker defeats his justification, and therefore the audience no longer knows it. He could, for example, create a undercutting defeater. Imagine that the speaker asks his friend how they know that the restaurant closes at 10pm, and when the dupe says that it says so on the website, the speaker could (falsely) say that this restaurant's website is unreliable or outdated. In all these cases, if knowledge is what we're talking about, then it's the case that the victim has been epistemically harmed in virtue of losing knowledge. If what we care about is having *justified* beliefs, but not knowledge, then the dupe is epistemically damaged in the latter case.

One account might have it that someone is epistemically damaged if their total number of beliefs is diminished. It's hard to imagine in what kind of case this will be a useful measure of damage, and so I'm not inclined to say much more about this.

The causing of a loss of beliefs might not count as damage, just as the loss of skin, let's say, might not count as damage, depending on what skin it is, or how it's removed. It could be that only the loss of (or losing justification for, or knowledge of, etc.) *important* beliefs is what matters. It's hard to give an account of what counts as *important* beliefs, and it's probably wrong to say that this "important" criterion is an epistemic one, and not a practical one. Perhaps one epistemic notion of importance is how central it is in one's web of belief. If one were to give a more sophisticated account of *epistemic importance* of any one belief, one might determine how many other beliefs would be affected by a loss of this belief, and take into consideration *their* importance as well. This could be a way to count this as an epistemic and not practical distinction. If, however, one were to say that a belief about one's own medication schedule were more important than one's beliefs about the Napoleonic Wars, this, I think, would not be a measure of epistemic damage, but practical damage. Of course, this will affect the blameworthiness of someone who lies about these topics, but I want to keep this distinct from the category of *epistemic* harm.

One last way we might conceive of epistemic damage is as harm to a person's rational capacity, which would inhibit their ability to *rationally* form or retain beliefs or obtain knowledge. If you are punched in the nose, and this causes sustained neurological harm such that your belief-forming and evidence-evaluating abilities are impaired, we may be inclined to count this as a kind of epistemic harm. Or, if the blow causes you to forget the last hour of events, *this* might count as epistemic damage. Perhaps this could be measured with an IQ or memory test. It's hard to know where this departs from practical harm, but it doesn't seem to me out of bounds to think that by impairing someone's capacity

to reason, you've harmed them, epistemically. I don't think that this is the most useful notion of epistemic damage, but it's at least plausible.

2.1.4 THREE CLARIFICATIONS

So far, I've talked about a person's being damaged because they are made to have *a* false belief, or their credence in *a* proposition has become more inaccurate. But this might not be the whole picture. In this section, I'll raise three different choices the proponent of an account of epistemic damage will have to make, or at least, consider. Much like determining in virtue of what it can be truthfully said that someone was epistemically damaged, the choices we make here will likely be context–sensitive, and the right choice will depend on the specific purpose to which we put epistemic damage.

Up until this point, I've talked about the epistemic damage a person suffers with respect to the accuracy of their credence in a proposition. In the paradigm case, one suffers epistemic damage when they become less confident in a true belief. But, if someone is damaged with respect to p, but benefitted with respect to q, then should we say that the person has been epistemically damaged? Should we just say that they were damaged with respect to p and benefitted with respect to q, or does it just make sense to say that they are no worse and no better off? Or, should we say that the damage with respect to p is incommensurable with the benefit to q? Or, if they've been damaged with respect to p and benefitted with respect to q and r, then should we say that they've been benefitted, overall. Or, does it depend on the relative *importance* of the beliefs at issue?

11

Here again, the answer will depend on the context in which the question is asked. If someone goes into cardiac arrest, and an EMT breaks that person's rib in performing life-saving CPR, we'd say that that person was overall benefitted. Nevertheless, it'd be odd to say that they didn't suffer any damage; indeed, they'd need to recover from the injury to their rib. The context, then, would determine whether it makes sense to talk about local or global damage.

A related question is whether we ought to think about damage at one time or over time. That is, how should we discuss the case where someone is damaged with respect to a particular belief, but, in the future, that they were damaged in this way caused some greater gain? To take another medical example: a child who receives immunizations suffers harm to their arm, but is, in the future, less likely to get diseases. Should we say, of this child, that they've been damaged? In some sense, it would be odd to answer "no" and at the same time administer a bandage. The better answer is to say "yes" but that overall, the child is better off, and perhaps second, that the small amount of damage was therefore justified. We can imagine that, in some circumstances, suffering damage with respect to one belief could cause a future benefit. Consider lies parents routinely tell their children: "You can do anything you set your mind to" and the like. The child is damaged with respect to *that* proposition, but let's just stipulate that believing it enables the student to study harder, and thereby, in the future, come to benefit epistemically. Or, we might include a small lie that teachers tell their students about a difficult concept or an author's view. Having a false belief about that concept might thereby make it easier to, in the future, gain a true belief (or many).

12

Last, consider when a speaker damages some, but benefits others, epistemically. Should we say, of that person, that they've done epistemic damage? Here, again, it seems that context will determine what the salient notion of damage is.

2.1.5 HOW ONE INFLICTS EPISTEMIC DAMAGE

I'll just say briefly the ways in which we might think that someone can inflict epistemic damage. Here's a non-exhaustive list: testimony; leaving misleading clues; omission, maybe; sending someone to a library where the non-fiction section is full of books that give false information; giving someone a concussion. The paradigm case of inflicting epistemic damage is an instance of lying where the liar says something that is false and where the dupe trusts the liar.

2.2 Expected epistemic damage

The first task of this paper was to come up with some method of *assessing* epistemic harm, which I've done by advancing an accuracy-based account of *epistemic damage*. The task of this section is to draw a line from epistemic harm to moral blameworthiness. I argue that when someone misleads, or lies, or leaves out misleading evidence, that person imposes an epistemic risk to their interlocutor (or, if they don't have someone in mind, *any* person who might come along). The epistemic risk is also a practical risk. The extent to which the speaker thinks they've put their interlocutor at epistemic risk, the *expected epistemic damage*, will determine, among other factors, their moral blameworthiness. Of course, expected epistemic damage does not fully determine moral blameworthiness. Lying about heart medication, for example, is a more blameworthy action than lying about the

13

weather, even if the expected epistemic damage is greater in the second instance. And, I don't claim that it's always morally wrong to inflict epistemic damage. My claims about the link between expected epistemic damage and moral blameworthiness are meant to be *ceteris paribus* claims.

2.2.1 EXPECTED EPISTEMIC DAMAGE AND MORAL BLAMEWORTHINESS

Above, I considered an instance of lying from the third-party perspective. But now consider an instance of lying from the *liar's* perspective: the liar has some low credence in a proposition that she nevertheless asserts to the dupe. The liar thinks the dupe will trust her, and therefore thinks that the dupe's credence in the proposition will increase. From the liar's perspective, her own credence maximizes expected accuracy. Therefore, according to the liar, anyone whose credence moves farther away from her own has become more inaccurate. While there is controversy about this, the general consensus in the literature is that you ought to be *immodest* about the accuracy of your own credences. This means that you should think that the expected accuracy of your own credences is greater than the expected accuracy of any other credences you might adopt.[7] The intuitive thought is that it is irrational to (i) think that there are different credences with a higher expected accuracy and (ii) not adopt them as your own.

Note that while epistemic damage depends on the truth-value of the proposition the liar asserts, *expected* epistemic damage does not. For example, if the liar mistakenly asserts

---

[7] See the literature on epistemic immodesty, e.g., Joyce, "Accuracy and Coherence," sec. 8; Horowitz, "Immoderately Rational."

a true proposition believing it false, and the dupe increases her credence in the assertion, the dupe does not suffer epistemic damage because she is epistemically better off. Nevertheless, the liar *expected* the dupe to suffer, so there is expected epistemic damage. Note also that the expected epistemic damage is not the same as the liar's prediction about the dupe's change in credence. Because the liar's statement does not vary according to her confidence—"It won't rain" sounds the same, we'll stipulate, no matter how confident the speaker is—we should expect the dupe's change in credence to be independent of the liar's credence. However, when the liar is *more* confident that what she asserts is false, she'll expect the dupe to become *more* epistemically worse off than in the case in which she is not as confident that what she asserts is false. Expected epistemic damage is a measure of the liar's estimate of how badly off the speaker will be, which is relative to whatever the liar thinks is the (maximally) rational credence to hold: her own. Expected epistemic damage is, of course, to be calculated using the speaker's degree of belief.

I claim that the greater the expected epistemic damage, the greater the liar's blameworthiness. The more confident the liar is that what she asserts is false, the more inaccurate she will think the dupe will become. The more inaccurate the liar thinks the dupe will become, the more likely the liar will think the dupe will not get what she wants. Once we adopt a degree-theoretic model of belief, it becomes clear that expected epistemic damage will vary according to the confidence the liar has that what she asserts is false. If we accept that epistemic damage correlates with practical harm, then it becomes clear that expected epistemic damage is a prediction of how badly off the liar thinks that her lie will make the dupe. In this way, a liar's blameworthiness is explained by the risk she imposes

on the dupe. Furthermore, this blameworthiness admits of degrees: the greater the liar's confidence that she asserts something false, the greater the risk she'll think she's imposing on the dupe, and, therefore, the greater her blameworthiness.

2.2.2 RISK

We can ask, separately, whether a risk is a harm and whether we are blameworthy for imposing risks on others. While most argue that imposing a risk doesn't in itself constitute a harm, some have argued that it can.[8] In the civil law, this question has garnered serious attention because only if there is a harm can there be a tort.[9] Put this question aside, though. What is uncontroversial is that we can be *blameworthy* for imposing risks on others, even if imposing a risk does not itself constitute a harm, and even if blameworthiness does not guarantee civil liability.

2.2.3 EXPECTED EPISTEMIC DAMAGE WITHOUT ACCURACY

Above, I gave an account of expected damage in terms of expected accuracy. But, I realize that not everyone will be on board with this. So, I want to sketch just how the notion of expected damage could work on a model without accuracy.

Say that what we care about isn't accuracy, but truth. That is, of some subset of the propositions (the important ones, let's say), we want to believe the true ones and disbelieve

---

[8] *See,* for arguments that mere exposure to risk can constitute a harm, Finkelstein, "Is Risk a Harm?"; Placani, "When the Risk of Harm Harms."

[9] That is, the courts and most scholars have rejected "risk-based liability" in favor of "damage-based liability." See, e.g., Porat and Stein, *Tort Liability under Uncertainty*; Robinson, "Risk, Causation, and Harm," 325 fn15.

the false ones. And we only either believe, disbelieve or refrain from belief. On this account, one suffers epistemic damage just in case one comes to believe a falsehood, or ceases to believe a true proposition. Expected damage, on this model, would be the liar's subjective likelihood that the damage would obtain. And, the greater the liar's subjective likelihood that the damage would obtain, the more likely she'll think that the dupe won't get what she wants. Thus, here again, the greater the expected epistemic damage the greater the moral blameworthiness, *ceteris paribus*.

2.2.4 ON WHAT THE EXPECTED DAMAGE DEPENDS

The speaker's assessment of her dupe's expected accuracy will depend on at least four things: 1) the speaker's credence; 2) the speaker's estimate of what the dupe's starting credence is; 3) the speaker's estimate of the extent to which the dupe trusts her; 4) the speaker's scoring rule. I'll address each, in reverse order.

Epistemic accuracy is determined by the agent's credence in a proposition, the truth value of that proposition, and the agent's scoring rule. While some scoring rules are more popular than others, one might have an idiosyncratic scoring rule, or, an agent could use different scoring rules. What kind of scoring rules the agent uses will determine the accuracy, and, therefore, expected accuracy, of any credence.

For a given communicative interaction, the more the audience trusts the speaker, the more she will increase her credence as a result of the speaker making an assertion. Note, this is *not* equivalent to the following: the more the audience trusts a speaker, the more she will increase her credence. Of course, we can trust people immensely but not increase our

17

credence (much) in light of what they say. For example: if your dearest friend tells you that a meter is longer than a yard, you wouldn't increase your already extremely high credence in that proposition. And, you likely would increase your credence in what time it is if you asked a stranger on the street. Yet this does not entail that you trust the stranger more than your friend. I mean here to say only, of a particular case, the more the audience trusts the speaker, the more they will increase their credence. In order for the speaker to have a view about the extent to which the dupe will be epistemically damaged, the speaker will have to think *something* about the extent to which the dupe will trust her.

One might here object, then, that this requires that in order for speaker to calculate expected epistemic damage they must know a lot about their audience. But, this is not the case. In some cases, one might have such knowledge. But, in many other cases, use of general information would be sufficient for an approximation. Imagine that the speaker is a doctor and the audience her patient. Without knowing much about the patient, the doctor may use general background knowledge about how trusting patients are of doctors, or female doctors, or doctors when they give bad prognoses. In addition, the audience can use background information about how trustworthy audiences are about particular kinds of assertions. It's plausible to think that people are more trusting of speakers who assert the kinds of propositions that are more likely to be true. For example, a speaker making an outlandish lie would be wrong to expect her audience, barring any particular knowledge about the audience, to trust her much.

Even if the speaker has a scoring rule and knows exactly the extent to which her audience would trust her, she won't be able to estimate the expected epistemic damage

18

without some inkling of her audience's starting credence. Expected epistemic damage is a measure, by the speaker's lights, of the distance between what the dupe's accuracy will be, once she updates on the speaker having made an assertion, and what it was before. But, in order to estimate where the audience will end up, the speaker has to have some idea of where she started. Again, in the absence of particularized knowledge about the audience, the speaker can use general, background knowledge. For example, speakers should estimate most people's credence in *snow is white* as very high, even without any particularized information about her audience. If they were they to lie to someone about the color of snow, only if they had reason to think that that person would really trust them should the expected epistemic be great.

Last, expected epistemic damage depends on the speaker's own credence in the proposition. Of course, what the speaker's credence is won't affect what the audience's credence is. Because the speaker's *statement* does not vary according to her confidence— "Snow is black" sounds the same no matter how confident the speaker is—we should expect the dupe's change in credence to be independent of the liar's credence. However, when the liar is *more* confident that what she asserts is false, she'll expect the dupe to become *more* epistemically worse off than in the case in which she is not as confident that what she asserts is false. Expected epistemic damage is a measure of the liar's estimate of how badly off the speaker will be, which is relative to whatever the liar thinks is the (maximally) rational credence to hold: her own. Thus, expected epistemic damage is proportional to the speaker's confidence that what she asserts is false. If we ground the blameworthiness of lying in expected epistemic damage, then we can say that the greater

the speaker's confidence that she asserts something false, the greater the risk she'll think

she's imposing on the dupe, and, therefore, the greater her blameworthiness.

2.3 Lying, risk and accuracy

2.3.1 INTRODUCTION

A large literature has yielded near unanimity on two necessary conditions on

lying.[10] One lies about *p* only if one:

(1)   Says that *p*
(2)   Believes that not-*p*.

Philosophers have discussed what satisfies the first necessary condition and what, if any,

further necessary conditions there are. But, absent from the literature is any serious

discussion of the second necessary condition—what I'll call the *belief requirement*. In

---

[10] *See,* e.g., Montaigne, *The Essays*, 20; Frege, "On Sense and Reference," 66; Demos, "Lying to Oneself," 588; Waismann, *The Principles of Linguistic Philosophy*, 294; Siegler, "Lying," 131; Mannison, "Lying and Lies," 133; Chisholm and Feehan, "The Intent to Deceive," 152; Davidson, "What Metaphors Mean," 42–43; Fried, *Right and Wrong*, 55; Coleman and Kay, "Prototype Semantics," 28; Kupfer, "The Moral Presumption against Lying," 104; Primoratz, "Lying and the 'Method of Ethics,'" 52 fn 152; Shibles, *Lying: A Critical Analysis*, 31; Wiles, "Lying: Its Inconstant Value," 275; Adler, "Lying, Deceiving, or Falsely Implicating," 437; Green, "Lying, Misleading, and Falsely Denying: How Moral Concepts Inform the Law of Perjury, Fraud, and False Statements," 159; Williams, *Truth and Truthfulness: An Essay in Genealogy*, 96; Meibauer, "Lying and Falsely Implicating," 1376; Faulkner, "What Is Wrong with Lying?," 536; Fallis, "What Is Lying?," 33; Fallis, "What Liars Can Tell Us about the Knowledge Norm of Practical Reasoning," 348; Fallis, 348; Sakama, Caminada, and Herzig, "A Logical Account of Lying," 300; van Ditmarsch et al., "On the Logic of Lying," 42; Fallis, "Lying as a Violation of Grice's First Maxim of Quality," 564; Saul, "Just Go Ahead and Lie"; Sorensen, "Lying with Conditionals," 820; Arico and Fallis, "Lies, Damned Lies, and Statistics," 791; Lackey, "Lies and Deception," 246; Stokke, "Lying and Asserting," 33; Webber, "Liar!," 652; Rees, "Better Lie!," 59; Dworkin, "Are These 10 Lies Justified?"; Fallis, "Are Bald-Faced Lies Deceptive after All?," 82; Keiser, "Bald-Faced Lies," 464; Leland, "Rational Responsibility and the Assertoric Character of Bald-Faced Lies," 550; Mahon and Zalta, "The Definition of Lying and Deception"; Michaelson, "The Lying Test," 481; Stokke, "Lying and Misleading in Discourse," 106. *Slight* deviations from this account are Isenberg, "Deontology and the Ethics of Lying"; Carson, "The Definition of Lying*"; Fallis, "Davidson Was Almost Right about Lying."

addition, although epistemologists have long recognized that people have attitudes of *degrees* of belief instead of (or in addition to) full belief, philosophers have not considered the possibility that the true requirement on lying concerns not the speaker's belief, but rather her degree-of-belief.[11]

I begin with a claim about what makes liars blameworthy, when they are. I claim that the correct explanation of why liars are blameworthy includes the liar's imposing a risk on the audience. This blameworthiness admits of degrees: the greater the liar's confidence that what she asserts is false, the greater the risk she'll think she's imposing on the dupe, and, therefore, the greater her blameworthiness. From this explanation of what makes liars blameworthy, I arrive at a dilemma: either the belief requirement is wrong, or *lying* isn't a particularly interesting concept. I then suggest a principled replacement to the belief requirement: the *worse-off requirement*.

---

[11] I'm not the first, however, to consider the role of degrees of belief in the context of lying. In response to a claim by Sorensen (2010) that knowledge-lies do not involve deception, Staffel (2011) and Fallis (ms) argue that a dupe is deceived by a liar when the dupe's confidence in a falsehood increases, even if the dupe does not thereby come to *believe* the falsehood. (Fallis, "Bayesians Don't Tell Knowledge-Lies (and Probably Nobody Else Does Either)"; Staffel, "Reply to Roy Sorensen, 'Knowledge-Lies.'") Chisholm and Feehan mention that the extent of deception depends on the degree of belief the dupe comes to have in the falsehood the liar asserts. (Chisholm and Feehan, "The Intent to Deceive.") While these authors offer important precedents for including degree-theoretic reasoning in accounts of lying by thinking about the extent to which the dupe is deceived, my project is to discuss the *speaker* using degrees of belief, and, importantly, to give an account of blameworthiness for lying that depends on the relationship between the speaker's degree of belief and her expectation about what the dupe's degree of belief will be. Thanks to an anonymous referee for urging me to clarify this point.

Marsili is, as far as I can tell, the sole person to argue for an account of lying that explicitly takes into consideration the role the speaker's degree of belief ought to play. (Marsili, "Lying as a Scalar Phenomenon: Insincerity Along the Certainty-Uncertainty Continuum.") He replaces the belief requirement with a requirement that the speaker must be more confident that the proposition she asserts is false, than true. I show why this suggestion is not correct.

2.3.2 DILEMMA

By lying, the speaker does something that she believes will cause the dupe to become less epistemically accurate, and, therefore, less likely to be able to satisfy her desires. I've argued that once we think in terms of degrees of belief, we can appropriately explain the blameworthiness of lying as the causing of expected epistemic damage, which depends, among other things, on the confidence the speaker has that what she asserts is false. But, once we so-ground the blameworthiness of lying in expected epistemic damage, the orthodox account of lying faces a dilemma: either the belief requirement is wrong, or *lying* isn't a particularly interesting concept. In what follows, I explain why the dilemma exists and suggest why impaling ourselves on the first horn isn't all that bad.

I've argued that blameworthiness for lying is best explained by the liar's imposition of risk on the dupe, which I analyze in terms of expected epistemic damage. If the belief requirement is correct, then there are nearby cases of non-lies in which the speaker is blameworthy for the same reason as the liar, albeit to a lesser degree, without lying. However, it's difficult to see why the belief requirement marks the distinction between lies and assertions that, by failing to meet the belief requirement, fall just short of lying. I'll propose, in the next section, that we can make the distinction more intuitive if we abandon the belief requirement and instead introduce a requirement involving degrees of belief.[12]

---

[12] Others have offered accounts of lying that generate other ideas for drawing a principled distinction between lies and blameworthy assertions that fall short of lying. Benton (2018), for example, argues that we ought to replace the belief requirement with a knowledge requirement. On his view, even if the knowledge requirement is not met, and thus one doesn't lie, one can be blameworthy if one "negligently" asserts a proposition, where one negligently asserts a proposition just in case "one doesn't consider whether one knows *p*." Benton, "Lying, Belief, and Knowledge," 128.

To get a sense of why the belief requirement is not natural or intuitive, imagine a case in which someone is not so confident that we would ascribe *belief* to her—say, she has .6 credence that it will rain. If the speaker nevertheless says that it won't rain, the speaker will think that her audience will be worse off, for trusting her. This is the case even if the speaker's confidence is not sufficient for belief, and therefore, according to the orthodox account, not sufficient for lying. But, we can explain what's wrong about this assertion in the same way as what's generally wrong about lying. The speaker's assertion that it won't rain when she has .6 credence that it will, and an assertion that it won't rain when she *believes* that it will differ only by the magnitude of risk imposition, and, therefore, the speaker's blameworthiness. If this is the case, it's hard to see why *lying* picks out any special phenomenon.

Imagine if we were to define a "theft" as an intentional taking of someone else's property, worth at least $100. One might balk at the $100 requirement and ask for a principled reason why a $40 taking shouldn't count as a theft, too. Perhaps there is a pragmatic explanation: perhaps a statute requires police to investigate all "thefts" but that it's inefficient to investigate takings of less than $100. So, defining "theft" in this way makes pragmatic sense. One might think that there is some similar, philosophically pragmatic reason for the belief requirement, but it's far from clear that there is a philosophical equivalent to the pragmatic classification of legal concepts, and, even if there were, what that justification would be in this case. If the belief requirement is correct, then lying does not pick out an especially interesting concept.

To address the first horn of the dilemma, if lying really does pick out an interesting concept, then the belief requirement must be wrong. One proposal is to replace the belief requirement with what Marsili (2014) calls the "comparative insincerity condition," according to which the speaker must be more confident that the proposition she asserts is false than true. However, this doesn't account for the damage lying does. Imagine a case in which the speaker is just slightly more confident that p is true than false—say she has .51 credence in p. And, further, imagine that the speaker knows that her audience is agnostic—that the audience has .5 credence in the proposition. From the speaker's perspective, the audience's credence is very close to her own, and, therefore, the speaker thinks the audience's credence nearly maximizes expected epistemic accuracy. When the speaker asserts that p, supposing the speaker believes the audience will trust her, she'll think that the audience will assume that she is far more confident than she is. Imagine that the audience moves from .5 credence to .8 credence. If the speaker's credence is .51, then the speaker will expect the audience to suffer epistemic damage. If the blameworthiness for lying is grounded in expected epistemic damage, then this speaker is blameworthy in exactly the same way as liars are, even if, according to both the orthodox position and Marsili's proposal, she hasn't lied.

2.3.3 THE WORSE-OFF REQUIREMENT

With the belief requirement, lies are not interestingly distinct from nearby cases. My proposal would save lying from being uninteresting. Using the simple, general

explanation I've given for why liars are blameworthy, it's possible to construct a principled

replacement to the belief requirement:

The worse-off requirement

> The expected epistemic damage to the audience, with respect to p, by the speaker's lights, conditional on the audience trusting her with respect to p, at all, is greater than 0.

That is, it's a necessary condition on lying that the speaker think that, if the dupe trusts her,

the dupe will be worse off, epistemically. I've deliberately made the requirement

conditional on the audience trusting the speaker; it puts off a decision about whether intent

to deceive is a necessary condition on lying. For those who argue that bald-faced lies

(assertions by a speaker who believes both that the proposition she asserts is false and that

her audience won't trust her) *are* lies, then the two above conditions might be sufficient for

lying.[13] Those who argue that intent to deceive is a necessary condition on lying would add

a further condition: that the speaker expect the dupe to trust her. My worse-off requirement

takes no position in this debate.

My view allows that some assertions that are not lies can be blameworthy (of

course), and I also allow that the blameworthiness of some of those assertions can be given

a closely related, structurally similar explanation using the notion of expected epistemic

damage. Where I draw a principled, natural boundary around lying, however, is when

damage is done with respect to p, the proposition asserted. I claim this is more intuitive

than where advocates of the belief requirement must draw the boundary.

---

[13] See the literature on bald-faced lies, e.g., Sorensen, "Bald-Faced Lies! Lying Without the Intent to Deceive"; Lackey, "Lies and Deception"; Fallis, "Are Bald-Faced Lies Deceptive after All?"

I'll note here two further features of the worse-off requirement. First, it is compatible with epistemic models that eschew belief entirely in favor of degrees of belief.[14] If the belief requirement is correct, it's not clear what lying is for a Bayesian without belief, for example. My replacement to the belief requirement provides an answer. Second, the worse-off requirement is consistent with including knowledge-lies as lies. According to Roy Sorensen, an "assertion that *p* is a knowledge-lie exactly if intended to prevent the addressee from knowing that *p* is untrue but is not intended to deceive the addressee into believing that *p*."[15] Sorensen claims that knowledge-lies do not involve deception, but Staffel 2011 and Fallis (ms) convincingly show that the interesting cases do, because even though the dupe does not come to *believe* the false assertion, she becomes more confident in it. The worse-off requirement is consistent with an account of lying according to which most, though perhaps not all, knowledge-lies count as lies.

I've argued that once we think in terms of degrees of belief and accept a simple explanation for the blameworthiness of lying, we arrive at a dilemma. To tackle that dilemma, I've suggested a principled alternative to the belief requirement—the worse-off requirement. Although nearly every philosopher who has given an account of lying has endorsed the belief requirement, my revision is not all that radical. The worse-off requirement recasts lying as involving an imposition of risk on the audience and better

---

[14] *See,* e.g., Jeffrey 1970, who writes that "our ordinary notion of *belief* is only vestigially present in the notion of degree of belief" Jeffrey, "Dracula Meets Wolfman: Acceptance vs. Partial Belief," 171–72.; or see Christensen 2004, who demotes the importance of binary belief for epistemic rationality. Christensen, *Putting Logic in Its Place*, 96–105.

[15] Sorensen, "Knowledge-Lies," 610.

captures what usually makes lying blameworthy: that the speaker expects the dupe to be worse off, for trusting her.

2.4 Objections

I've argued that if lying is an interesting concept, then the belief requirement is wrong. I've replaced the belief requirement with the *worse-off* requirement, and also argued that the more confident the speaker is that what they assert is false, the more blameworthy they are for lying. In doing so, I've favored degrees-of-belief over outright belief. One might take this to entail, or suggest, that lying, then, is a graded phenomenon. I won't argue that this is *not* the case, but I want to resist the inference to it from anything I've said, above. Just because belief is graded, and belief (in some way) is a constitutive element of lying, this does not mean that lying itself is graded. For example, say that theft is the intentional taking of another's property. Thus, taking some amount property is a constitutive element of theft. Amounts of property come in degrees: $1, $5, etc. It doesn't follow that a theft of $100 is more of a theft than that of $1. Granted, it is a theft of more money, and the thief might be more blameworthy and subject to increased penalties, but it doesn't entail that one is more of a theft than the other.

Another objection might be that, on my view, with one statement one can lie to one subset of an audience but not the other. Imagine a case in which a speaker has two audience members: one gullible, the other incredulous, and the speaker knows this. The speaker says something she believes is false. On my view, there will be expected epistemic damage for one but not the other. And, if the blameworthiness for lying depends on the expected

27

epistemic damage, then it seems to follow that the speaker is both blameworthy and not blameworthy for lying. Both of these consequences might seem odd, and thus might pose a problem for my view.

If this objection poses a problem for my view, though, it is a problem for lying generally, at least for those accounts of lying according to which the speaker must intend to deceive her audience. Take any case in which a speaker knows that her audience members have differing levels of credulity. If lying requires intent to deceive, then the speaker, in making one assertion, will have lied and not lied. And this might seem like a bad result. But this does not seem troubling to me. If lying requires an audience, then it seems that, in saying one thing, a speaker can lie to some and not to others. And this is the same on my view: in saying one thing, one can lie to some but not to others. One can be blameworthy for lying because of the harm to one person, even if another is unharmed.

As a brief aside, this is not unique to lying. Consider someone who, in a body of water, flips over a raft on which two people are sitting: one who can swim, and one who can't. With one action, the person can both harm one person and not harm another. She can be blameworthy for what she does to one person, and not to another. Of course, she might be all-things-considered blameworthy for flipping the raft, but this is true in the lying case as well.

If, as I say, epistemic damage correlates with practical harm, and we can say that expected epistemic damage is a prediction of how badly off the liar thinks that her lie will make the dupe, then one might object that we should just skip the expected *epistemic*

28

damage altogether and think only about the practical consequences.[16] That is, if expected epistemic damage is useful only as a proxy for practical damage, what's the point of it?

Basically, I want the notion of epistemic damage to be used in many different kinds of circumstances. One communicative interaction, causing some fixed amount of epistemic damage, may lead to several different inferences about blame. Or, whether the epistemic damage inflicted in a particular case causes actual harm will depend on the circumstances. Other times, we might not know the practical consequences of a communicative interaction, and measuring the epistemic damage is a good stand-in. Take, for example, blood loss. Presumably, we care about blood loss only insofar as it affects our health. Why, then, should EMT's say anything about blood loss of a patient, and not instead just say what the harms are, or what the remedy ought to be? For one thing, whether the blood loss is salient will depend on the circumstances. The prescribed treatment might not depend only on the amount of blood loss, but also the desires of the patient, or the status of the patient's kidneys, or other factors. Maybe sometimes blood loss is good: when someone donates, for example. But the amount of blood a patient has lost is an important piece of data that medical staff use in treating patients.

One might object that my replacement to the belief requirement is consistent with the classification of some assertions as lies that the traditional account would exclude. And, one might object that some assertions that might have been classified as lies on the traditional account, no longer would. My inclination is to bite the bullet on these kinds of

---

[16] Thanks to Kenny Easwaran for raising this concern.

cases. In general, there are cases in which one asserts a proposition one believes is false in which the worse-off requirement is not met, and there are cases in which one asserts a proposition one believes is true and the worse-off requirement *is* met. If these cases seem to lead to unintuitive results, my strongest response is that's because the orthodox account of lying is either uninteresting or else has the wrong necessary condition. A weaker response is that I'm merely providing a necessary condition, and that in those cases where the worse-off condition is met, and yet it does not seem like a lie, anyone inclined could articulate further necessary conditions to rule out the case as an instance of lying. For now, let me address, at some length, one kind of case one might make.

> The speaker, A, has .6 credence in *p*, and knows that the audience, B, has .9 credence in *p*. A has good reason to believe that, because of the nature of B's evidence, B won't decrease her credence much below .6, even if A says that *p* is false. In a charitable effort to bring B's credence closer to A's own, and knowing that it won't fall much below .6, A says that *p* is false.

Assuming that .6 credence in *p* is sufficient for belief, most, I think, still wouldn't count this as a lie. That the liar must intend to deceive the audience is a view shared by most who've written on lying, though not all.[17] In the case as described, A's intent is to make B's credence *more* accurate. And, A has good reason to think that, given the unusual and somewhat artificial nature of the thought experiment, by asserting that *p*, this will be the case. Thus, A does not intend to deceive B. So, assuming the plausible intent-to-deceive requirement, this is not an instance of lying. At worst, this case shows that the worse-off requirement works only on views according to which intent to deceive is a necessary

---

[17] *See*, e.g., Carson, "The Definition of Lying*"; Sorensen, "Bald-Faced Lies! Lying Without the Intent to Deceive"; Fallis, "What Is Lying?"

condition on lying. However, I've intentionally made the worse-off requirement consistent with accounts of lying that do not require intent to deceive, so it'd be better for my purposes to not make this concession. So, let me add a further argument.

The case as described ought to be counted as a lie only if we adhere to the belief requirement and don't require intent to deceive. The wide array of accounts of lying that include the belief requirement, I think, fail to properly appreciate the normative properties of lying. I've argued that a good account of lying ought to capture why liars are blameworthy. The belief requirement leaves out a great number of these kinds of cases. The worse-off requirement explains, better than the belief requirement does, what makes liars blameworthy—namely, that the liar expects her audience to be epistemically worse off. If this unusual case seems intuitively to count as a lie, then I think it's a justifiable cost to my proposal that it excludes this case in exchange for the great gain of many cases that share the interesting features of lies, but, on the traditional view, are excluded.

Matthew Benton (2018) proposes three counterexamples to my revised account of lying.[18] However, the purported counterexamples fail. Essentially, what Benton shows is that, with my revision, some cases that would have counted as lies by the traditional account, don't, and others that would not have, do. Of course, I count this as a virtue of my revision. Benton's strategy is successful only if we are already committed to the traditional account, which is the very thing I seek to challenge.

I'll now discuss each of the purported counterexamples.

---

[18] Benton, "Lying, Accuracy and Credence."

Case 1: A's credence in p is 1, B's in p is 0 (and 1 in ~p), and A knows this. A asserts to B that ~p.[19]

Benton correctly notes that this case fails WORSE-OFF: Because B is already maximally confident that ~p, B cannot be epistemically worse-off, by A's lights. Therefore on my account it is not a lie. It is purportedly a counterexample because "A is maximally certain that what A asserts is false, so A's assertion is intuitively a lie."[20] But this explanation of the case as a counterexample fails for two reasons.

First, the only justification Benton gives for Case 1 to serve as a counterexample is that it meets the belief requirement (or, one might say, because of A's *certainty*, exceeds it). The whole point of my argument is to challenge the belief requirement. Putting forward a case whose only merit is that it meets the belief requirement, as a counterexample against an account that rejects the belief requirement, is question-begging.

Second, even if the case meets the belief requirement, it fails a requirement that most, but not all, believe to be a necessary condition on lying: that the speaker intends to deceive her audience. In the case as presented, A can't both know that B's credence in p is 0 and intend to deceive her about p: Necessarily, when one has 0 credence in a position, no new information can change that. Because A knows that nothing she says about p will change B's mind, she can't intend to deceive B.[21] Thus, the case is not a lie according to

---

[19] Benton, 196.

[20] Benton, 196.

[21] One possible rejoinder is that A might succeed in deceiving B about whether *she*, A, believes that p. That is, even if A won't change B's confidence with respect to whether it will rain tomorrow, she will be able to trick B into thinking that *A believes it will rain*. I take it, however, that of those who think that the intention to deceive is a necessary condition on lying, the thing about which one must intend

most people. For this reason Case 1 is an odd choice if it is meant to be an obvious example of lying. It's also an odd case because there are few propositions in which (rational) people have zero credence, and even fewer where their interlocutor knows this.

> **Case 2**: A's credence in p is 1, B's in p is 0.1 (and 0.9 in ~p), and A knows this. A asserts to B that ~p.[22]

Benton correctly notes that Case 2 passes WORSE-OFF, and therefore on my account could be a lie. I say "could be" because WORSE-OFF is only a necessary condition. Assume it is a lie, though.

That one could count Case 2 but not Case 1 as a lie is problematic to Benton. In both cases, Benton writes, "…A has the same high confidence that the asserted p is false, and intuitively both cases are lies if either of them is."[23] Put another way, because in both cases A believes that p is false, whether A lies shouldn't vary between the cases.

However, Benton is wrong to think that simply because in the two cases the speaker has the same high confidence in the falsity of the proposition she asserts the speaker is either lying in both cases or not lying in both cases. You don't have to buy my revisionary account to see that Benton is incorrect here. Anyone who believes that intent to deceive is a necessary condition on lying should immediately realize that two cases, keeping the speaker's confidence constant, can differ in whether they are lies. Take the following pair of cases:

---

to deceive their audience is the truth of proposition asserted, not (only) the speaker's attitude toward the proposition asserted.

[22] Benton, "Lying, Accuracy and Credence," 196.

[23] Benton, 197.

Case 4: A speaker wearing a white sweater stands in front of an audience and says "I am wearing a blue sweater."

Case 5: Same as Case 4, except the audience is blind, and the speaker knows this.

Case 4 is what has been come to be called a *bald-faced lie*.[24] Clearly, there is no intent to deceive the audience. Assuming, with the majority of scholars, that intent to deceive is a necessary condition on lying, Case 4 is not a case of lying. And, clearly, Case 5 *is* a case of lying. Thus, two cases where the speaker has the same (extremely low) confidence in what she asserts can differ as to whether the speaker lies. That the WORSE-OFF view treats Cases 1 and 2 differently is no evidence against my revision.

Case 3[25]: A's credence in p is 0.8, B's is 0.8 in p, and A knows this. A asserts to B that p; and A expects this to raise, even slightly, B's credence in p.[26]

Benton's diagnosis of this case is that, on my account, A lies by "telling the truth," which does, it may seem, pose a problem for my account. But, I have a few ways to respond.

First, it's not clear that the case as described would involve an intent to deceive, and thus, would not count as a lie (on most views). In Case 3, the speaker knows that she and the audience have the same, high credence. Under these circumstances, a reasonable person might not think that her assertion would cause any change in the speaker's credence.

---

24 Whether or not a bald-faced lie is a lie is a matter of some contention. See, e.g., Sorensen, "Bald-Faced Lies! Lying Without the Intent to Deceive"; Lackey, "Lies and Deception"; Stokke, "Lying and Asserting"; Fallis, "Are Bald-Faced Lies Deceptive after All?"; Leland, "Rational Responsibility and the Assertoric Character of Bald-Faced Lies"; Keiser, "Bald-Faced Lies."

25 (The numbering of the cases is intentional, to match those that appear in Benton 2018.)

26 Benton, "Lying, Accuracy and Credence," 197.

Keep in mind, while the speaker knows what both of their credences are, and knows they are the same, the audience does not. B doesn't know either (i) what A's credence is, or (ii) that A knows what B's credence is. It's plausible that the audience, B, could think that the speaker, A, would assert that *p* even if she had less than .8 credence in *p*. B will likely think that A has *some* high credence in *p*, but B could think that A could speak sincerely and yet have a *lower* credence than B does in *p*. If this is the case, then B will not count A's having said that *p* as evidence for *p*, and thus, will not increase her credence. A, knowing all of this, then, will not think that B will be epistemically damaged. Thus, while in this case WORSE-OFF is met, an extremely plausible, widely-held requirement, intent-to-deceive, is not, and so, the speaker has not lied.

Second, even accepting that intent to deceive is a necessary condition on lying, one might be concerned that WORSE-OFF allows that someone can lie by asserting a proposition that they have high credence in—something they believe.

Though these cases may occur, they will be rare, and, when they do occur, the speaker would be less blameworthy than in the ordinary case of lying because the expected epistemic damage will be low. And, that WORSE-OFF is consistent with these kinds of cases makes sense given the reason lying is wrong. If we agree that lying is wrong, when it is, because it causes the audience to be worse-off, epistemically, then those case in which someone is made worse-off, epistemically, are good candidates for lies. This is the case *even if* the speaker asserts a proposition she has high confidence in.

Third, I give WORSE-OFF as a necessary condition; it's consistent to add a further necessary condition to assuage those particularly troubled by Case 3. For example:

One lies about p *only if one doesn't believe that* p.

One lies about p *only if one doesn't have high credence (or a credence of at least x) in* p.

I think the burden is on the philosopher who would want to restrict the class of lies to show why such a restriction is warranted, but I leave that project to any interested party.

Benton believes that the above counterexamples show that my revision to the traditional account of lying is a bad revision. I've argued, above, that all three purported counterexamples fail. But, even if they *do* show that there are problems with my account, they do not show that the traditional account is superior. Benton does not address my main argument, which is that the belief requirement gets wrong a whole class of cases and does not accord with the best explanation for what makes liars blameworthy. The worse-off requirement does it better.

3 MORAL MARKET DESIGN[27]

3.1 Introduction

People have legal entitlements to act in ways that others regard as morally impermissible. When they do, others have a defeasible reason to stop them. Circumstances will dictate whether others should try to inhibit such behavior, and, if so, the best method: one might convince such people that what they are doing is wrong; one might explain that others will dislike them if they persist; one might ask them nicely, or threaten them; or, one could pay them.

I argue that we ought to pay people to stop behaving in ways we believe are immoral but toward which they are morally indifferent. I argue for the creation of a *moral market*, where people buy and sell the cessation of such behavior. Properly regulated, this market would have two significant upsides, noted here, and few downsides. First, it would allow people to trade in a way that would make them better off. Second, it would bring about a moral improvement.

Of course, this is a highly theoretical, academic proposal. Before designing a new market, experts outside of philosophy need to weigh in.[28] I aim merely to draw attention to the inefficient allocation of certain entitlements and to offer a framework to address this problem.

---

[27] A version of this chapter appeared as Krauss, Sam Fox, "Moral Market Design," *Kansas Journal of Law & Public Policy* 28(2) (2019): 426–455, of which I am the sole author.

[28] *See* Ballantyne, "Epistemic Trespassing."

This chapter proceeds as follows: first, I use the case of the ethical vegetarian to motivate the proposal and articulate the positive argument for creating a moral market; second, I refine the proposal and set limits to the market; third, I consider several objections; fourth, I offer two additional proposals in the same spirit, which those initially unconvinced of the desirability of the moral market may find more palatable.

3.2 Moral Indifference and Allocative Inefficiency

Moral disagreement is ubiquitous, but it is nevertheless striking that we are surrounded by people who believe that morality requires them to act in ways incompatible with what *we* believe morality requires. There is a large literature on the problem of disagreement and what to do about it.[29] Chiefly, however, this article considers cases of disagreement in which one party is morally indifferent toward their own behavior while the other party believes that they act impermissibly. *This* kind of moral disagreement allows for both increased economic efficiency and moral progress, so far not discussed.[30]

Moral disagreement may be (generally) intractable, and therefore so too the actions that follow from disagreeing parties. But when the actions stem from non-moral reasons,

---

[29] *See generally* Enoch, "How Is Moral Disagreement a Problem for Realism?" (Arguing that moral disagreement does not pose a problem for moral realism); Audi, "Intuition, Inference, and Rational Disagreement in Ethics." (Arguing that intuitionism can allow for rational disagreement); Elga, "Reflection and Disagreement."(arguing for a broadly conciliationist view); McGrath, "Moral Disagreement and Moral Expertise."; Vavova, "Moral Disagreement and Moral Skepticism." (Discussing the extent to which moral disagreement gives cause for skepticism); Schoenfield, "Permission to Believe." (Defending a permissive view of epistemic rationality and discussing its connections to disagreement); Feldman, "Reasonable Religious Disagreements." (Arguing for a moderate skepticism in light of religious disagreement).

[30] One notable exception is Ord, "Moral Trade." (Arguing that differences in two parties' moral views allows for trades that constitute moral and prudential improvements over the status quo).

the actor can be deterred with financial incentives or in-kind payments. And, while people should not allow financial incentives to outweigh moral reasons, financial incentives often do outweigh non-moral reasons—innocuously so. One ought not, say, accept money to paint a racial slur on a building, but painters earn a living painting buildings they otherwise have no reason to. To motivate the proposal for the moral market, consider the case of the ethical vegetarian and the omnivore.

3.2.1 THE VEGETARIAN AND THE OMNIVORE

Ethical vegetarians believe that eating meat is morally impermissible.[31] Omnivores, (at least, those who are not akratic), believe that eating meat is morally permissible but not obligatory. That is, while omnivores believe that there's nothing wrong with eating meat, they do not believe that they *ought to* eat meat, just as they do not believe they *ought to* go to a Yankees game, or wear a blue shirt.[32] One can say that omnivores are *morally indifferent* toward eating meat and watching baseball and wearing blue shirts. Of course, there are non-moral reasons to eat meat, watch baseball, and wear blue. To give up those activities might come at a cost, but not a moral cost.

However much vegetarians detest meat-eating, they cannot just force meat-eaters to stop eating meat: people have a legal entitlement to do so. However, there are some

---

[31] Granted, it's a bit more complicated than this. On the most plausible consequentialist views, it's not the *eating* of meat that is morally impermissible—it's killing (and/or causing suffering) that's wrong. I'd like to put this aside and let "eating meat" stand in for whatever it is that ethical vegetarians of any stripe think that omnivores do wrong.

[32] With some exceptions. *See, e.g.*, Meyers, "Why It Is Morally Good to Eat (Certain Kinds of) Meat: The Case for Entomophagy."; Bruckner, "Strict Vegetarianism Is Immoral." (Arguing that under certain conditions vegetarianism is immoral).

meat-eaters who, for a price, would become vegetarians for a period of time.[33] And, presumably, there are some people willing to pay that price.[34] If one owns the entitlement to eat meat, then one can sell that entitlement to the vegetarian. In effect, the vegetarian is contractually binding the meat-eater to become a vegetarian.

Entitlement is a notoriously elusive concept.[35] Let *entitlement* mean the legal ability to engage in a behavior or activity. Ignore the distinction between conduct to which people have an enumerated right—say, to vote—and conduct that no law currently prohibits, but easily could—say, the right to park on a certain street. Thinking about the ability to eat meat as an entitlement, and arguing that, in some cases, it ought to be for sale does not entail that all entitlements are, or ought to be, alienable or salable.[36] One cannot, for example, give or sell oneself into slavery.

It might seem odd to think about buying and selling the entitlement to eat meat. After all, people are legally permitted to eat as much meat as they want, so, one might say, acquiring someone else's entitlement does not get them anything. But that is wrong. While in some cases of transfer of ownership one acquires a new good, or the ability to engage in

---

[33] If the reader finds this is implausible, consider the many boring, painful, risky things people do for money, and that being vegetarian is far easier than many (if not most) of them.

[34] If the reader finds *this* is implausible, consider: i) that the price some people would be willing to accept to become vegetarian might be quite low; ii) that (some) people who donate to People for the Ethical Treatment of Animals, for example, do so for this very reason; and iii) that committed vegetarians spend money on far less worthwhile things: in 2016 the average American spent $388 on shoes. U.S. Bureau of Labor Statistics, "Consumer Expenditure Survey."

[35] *See* Ayres and Balkin, "Legal Entitlements as Auctions," 704.

[36] *See* Dagan and Fisher, "Rights for Sale," 106–124; Rose-Ackerman, "Inalienability and the Theory of Property Rights."

a behavior one otherwise could not, here the buyer pays the seller to *lose* the legal ability to perform some behavior. In the meat-eating case, the vegetarian buys the entitlement from the omnivore not because the vegetarian wants to eat meat, but precisely because they want the omnivore not to.

A trade on the moral market would make both parties by their own estimation better off: the omnivore values the entitlement to eat meat less than the vegetarian values the omnivore becoming a vegetarian. Without this trade, some people to whom the entitlement was granted will not derive the full value of the entitlement.[37] That is, if the entitlement to eat meat would be worth $100 on the moral market, but someone values their entitlement at only $50, then they are effectively out $50.

The moral market is not necessarily shielded from regulation. While there is a presumption in favor of voluntary transactions, this presumption can be overcome in the case of market failure, in which case the state can regulate or prohibit trade.[38] The argument in favor of trades on the moral market is the same as the argument for trade, generally, with the added benefit that it would make things *morally* better, as well.[39] Given certain restrictions, trades on the moral market can withstand the objections that are made to limit

---

[37] *See* Dagan and Fisher, "Rights for Sale," 96.

[38] For discussion of the presumption of free trade, *see, e.g.*, Rose-Ackerman, "Inalienability and the Theory of Property Rights," 932.; *see generally* Sunstein, "Incommensurability and Valuation in Law." (Arguing that the government should respect individual autonomy in valuation); Levmore, "Voting with Intensity." (Noting the common justification that voluntary trades increases utility). For discussion of government *regulation* of trade in response to market failure, *see, e.g.*, Cooter, "Market Affirmative Action," 134. For discussion of government *prohibition* of trade in response to market failure, *see, e.g.*, Rose-Ackerman, "Inalienability and the Theory of Property Rights," 933.

[39] For a recent, robust defense of the market, *see generally* Oman, *The Dignity of Commerce*. (Arguing that well-functioning markets are essential to liberal society).

trade in other goods and services, like organs and sex, as well as objections unique to this market.

3.2.2 PAYING FOR BEHAVIOR, AND TYPE-1 CASES

In the vegetarianism case, the buyer wants the seller to cease immoral behavior. One can also imagine cases in which the buyer pays the seller to *do* something rather than to *stop* doing something. In particular, one can imagine cases where the buyer believes that the seller is failing to engage in morally obligatory behavior, and the seller believes that they are failing to engage in morally permissible, but not obligatory, behavior. Imagine, for example, that environmentalists believe that it is morally obligatory to recycle; others believe that it is morally permissible, but not obligatory, to recycle. On the moral market, the environmentalist could pay the non-recyclers to recycle.

These two kinds of cases, paying someone to cease immoral behavior and paying someone to engage in morally obligatory behavior, though structurally distinct, are not morally distinct. Assume that to fail to engage in morally obligatory behavior just is to act immorally. With this taxonomy in mind, note that the two cases share the following structure, which I'll call a "Type-1" case:

Type-1 Case   A believes B's φ-ing is impermissible. B believes their φ-ing is permissible but not obligatory.

Given some assumptions, when the case has the above structure there is reason to think that A ought to pay B to stop φ-ing.

Not only would trades on the moral market increase economic efficiency, but such exchanges would make things morally better, at least from the buyer's perspective, and

42

morally worse according to neither. Consider the perspectives of a buyer and seller, respectively, contemplating such a transaction. The buyer believes that the seller is engaging in immoral behavior, say, eating meat, and after the trade, would no longer. Assuming that the buyer does not give up anything of comparable moral value by spending the money, from their perspective this is a moral improvement.[40] The seller believes that they are engaging in morally permissible, but not obligatory, behavior–eating meat, and after the trade, would no longer. From the seller's perspective, the trade is morally neutral. Note two features of the views of the parties: first, from the perspective of the buyer and seller in aggregate, the trade constitutes a moral improvement; second, neither party thinks that the trade constitutes a moral loss. Granted, that neither party believes things are morally worse does not mean that they are not. I'll address this point in Section 3.3.4.

If a voluntary transaction can make both buyer and seller better off, from their own points of view, and can produce, from their aggregate perspectives, a net moral improvement, and neither party thinks there was a moral loss, then barring significant market failure it seems that people ought to engage in these kinds of trades. In Section 3.3 I'll consider limitations to trades on the moral market and refine the proposal to take them into account. In Section 3.4 I'll consider and respond to several objections. However, none of the objections are sufficient to undermine this central piece of the argument: given some

---

[40] For a discussion of sacrifices of comparable moral importance, *see* Singer, "Famine, Affluence, and Morality," 231. (Arguing that we ought to help others unless doing so would require us to sacrifice something of moral importance comparable to the good we would do by such a sacrifice).

assumptions, if Adam believes that Bethany is acting immorally, and Bethany is morally indifferent to her own behavior, then Adam should pay Bethany to cease her behavior.

3.3 Refining the Proposal

3.3.1 DIFFERENT KINDS OF PERMISSIBLE BEHAVIOR

I noted that the cases I described had the following, Type-1 structure:

Type-1 Case   A believes B's φ-ing is impermissible. B believes their φ-ing is permissible but not obligatory.

But this case is underspecified: people are not morally indifferent to all behaviors they believe are permissible but not obligatory. Consistent with the seller believing that φ-ing is permissible but not obligatory, there are several different ways to classify the behavior: it could be morally neutral, a morally permissible moral mistake, suberogatory, or supererogatory.

An action is *morally neutral* just in case either there are no moral reasons that weigh in favor or against that action, or the reasons for and against balance each other out.[41] People ought to be morally indifferent toward morally neutral actions. The vegetarian example is just this kind of case: omnivores are morally indifferent toward meat-eating. Paying someone to cease behavior that they regard as morally neutral was the impetus for this project. Both parties benefit, and, in aggregate, from their perspectives, the trade

---

[41] Others have used *morally neutral* in a similar way. *See, e.g.*, Liberto, "Denying the Suberogatory," 396. Pace, "The Epistemic Value of Moral Considerations," 266. Nadelhoffer, "The Butler Problem Revisited," 277. Broome and Morton, "The Value of a Person," 179. Driver, "The Suberogatory," 294. Dale Dorsey discusses the closely related concept, *amorality*. *See* Dorsey, "Amorality." (Arguing that, in addition to obligatory, permissible, impermissible, supererogatory, and suberogatory actions, there are also *amoral* actions, which lack moral status).

constitutes a moral improvement without either party thinking that things are morally worse.

But not all morally permissible actions are morally neutral. Supererogatory and suberogatory actions, and morally permissible moral mistakes, though morally permissible, are not morally neutral.[42] People have moral reasons to save others from burning buildings, and to change seats so a couple can sit together, though it is usually considered morally permissible to do neither.[43]

Even if the case involves paying people to do things that the recipient believes they ought to do, there are reasons to be wary, especially in the moral case.[44] There is reason to be concerned about paying someone to engage in behavior they believe they have moral reason not to, or paying someone to desist from behavior they believe they have moral reason to engage in. I exclude from consideration those cases in which the seller would be paid to do things they believe they have moral reason not to do, or to cease behavior they believe they have moral reason to do.

---

[42] On the supererogatory, *see* generally Urmson, "Saints and Heroes." (Arguing for the classification of some action as supererogatory, although, without using that name); Horton, "The All or Nothing Problem." (Describing and offering a solution to a dilemma that arises for supererogation). On the suberogatory, *see* Capes, "Blameworthiness Without Wrongdoing." (Arguing that moral permissibility does not entail lack of blameworthiness); Driver, "The Suberogatory." (Calling attention to this class of actions). On morally permissible moral mistakes, *see* Harman, "Morally Permissible Moral Mistakes." (Arguing for a new classification of actions, *morally permissible moral mistakes*, that we ought not do, for moral reasons, but doing them is not morally wrong).

[43] *See* Driver, "The Suberogatory," 291.

[44] *See* Bowles, *The Moral Economy*. (Arguing that we ought not incentivize with money conduct people know they morally ought to do).

3.3.2 TYPE-2 CASES

So far, discussion has been limited to circumstances where the seller regards the behavior they engage in as permissible but not obligatory. But the reader need not limit themself to these cases: imagine cases in which the buyer pays the seller to cease behavior that the seller believes is morally *obligatory*, or, equivalently, to engage in behavior that the seller believes is morally impermissible. Thus, *Type-2* cases:

Type-2 Case   A believes B's φ-ing is impermissible. B believes their φ-ing is obligatory.

What should the reader make of these cases? Imagine adherents of two different religions, each of which requires that all people belong only to it. Adam believes that Bethany's membership in Religion B (and lack of membership in Religion A) is morally impermissible. Bethany believes that her membership in Religion B (and lack of membership in Religion A) is morally obligatory. Stipulate that there is some amount such that Adam is willing to pay so that Bethany would be willing to convert to Religion A. While *some* people might be unwilling, at any price, to convert, this is not the case for everyone, even if they might feel torn, or, after the fact, guilty.

Many who were persuaded of the desirability of the moral market to handle at least some Type-1 cases will balk at Type-2 cases. Sure, they might say, pay someone who does not care to switch from Big Macs to veggie burgers, but to pay someone violate their moral beliefs is too much!

The proposal, then, is strictly limited to Type-1 cases, and specifically when the seller believes that they are engaging in morally permissible, but not morally *good*, behavior. Leave it an open question whether one ought to pay people to cease behavior that

46

they regard as morally impermissible—that is, whether one ought to pay others to stop acting akratically. For the remainder of the article, then, I will consider objections only to the limited positive proposal that follows: that one ought to pay others to cease behavior they believe is impermissible and toward which the sellers are morally indifferent.

3.3.3 LEGALITY

Only *legal* entitlements are traded on the moral market. Two kinds of transactions are excluded from consideration: paying people to cease immoral, illegal activity, and paying people to engage in morally obligatory, illegal activity. People do not have a legal entitlement to engage in illegal activity, and so there is no entitlement to sell. For example, on the moral market one cannot pay someone else *not* to use illegal drugs. No one has a legal entitlement to use drugs, so, at least on this framework, they have nothing to sell.[45] Also excluded from the proposal is paying people to engage in illegal activity that the buyer believes is morally obligatory. For example, on the moral market one cannot pay someone else to, say, remove animals from a research facility.[46]

---

[45] This is not to say, however, that private citizens have no reason to try to get others to obey the law, and it's not to say that paying people to obey the law is a bad idea. But I don't endorse this idea, and the moral market, as I propose it, cannot handle these kinds of cases.

[46] Here again, this is not to say that private citizens never have a reason to try to get others to break the law, and it's not to say that paying people to break the law is always a bad idea. In some exceptional circumstances, perhaps in an unjust society or under an oppressive regime, it may be warranted. But I don't here endorse this idea, and it is not part of the proposal for the moral market.

3.3.4 MORAL PROGRESS

In addition to increased allocative efficiency, the moral market will bring about moral progress, according to many views of what that could mean.

The proposal endorses trades on the moral market only in a subset of Type-1 cases: those in which the seller is morally indifferent toward their own behavior. The vegetarian and the omnivore represent a paradigm case. The trade constitutes a moral setback according to neither party, and a moral improvement if readers consider their views in aggregate.[47] If, in aggregate, the parties believe that things are morally better, and no party to the transaction believes that things are morally worse, then there is some reason to think that things are, actually, morally better off. This is why trades on the moral market would lead to moral progress.

There are two kinds of objections here. The first is that people's moral beliefs are not reliable, and thus, one cannot infer from the beliefs of the parties that the trade constituted a moral improvement to a claim that it did. The second is that, even if people's moral beliefs are, in general, reliable, trades on the moral market will be overrepresented by those with unreliable moral beliefs, and therefore a market in entitlements will not lead to moral progress.

If the objection to this argument is that people, generally speaking, have unreliable moral beliefs, then the objection runs far deeper than a mere opposition to the moral

---

[47] *See* Kornhauser and Sager, "Unpacking the Court." (Applying concepts of aggregation to judicial decision-making). *See also* List, "The Theory of Judgment Aggregation." (Providing an introduction to the philosophy of judgment aggregation); List and Polak, "Introduction to Judgment Aggregation." (Providing a technical introduction to judgment aggregation.)

market.[48] If people have unreliable moral beliefs, and therefore should not make trades informed by them, it seems also that we have reason to doubt whether they ought to act on them. Granted, there is reason to take this objection seriously, but this objection is not unique to the moral market proposal; and if it is a good objection, it raises serious doubts about whether people have reason to act in accordance with their moral beliefs.

But if the objection to this argument is, however, not that people are *generally* unreliable in their moral beliefs, but that unreliable believers would be overrepresented in the market, then there might be cause for concern. Perhaps the wealthy will participate more in the market, and the wealthy are more likely to have mistaken moral beliefs. But why should one think that the wealthy have comparably unreliable moral beliefs? One might say that this is because the wealthy tend to act immorally. Or, one might say that the wealthy necessarily act immorally—by not giving away their money and thereby becoming not-wealthy. However, this explanation is orthogonal to the original objection: even if the wealthy *act* immorally, this does not mean that their moral beliefs are unreliable. It may mean, instead, that the wealthy act akratically.

In addition, even if wealthy *buyers* are unreliable, the transaction would also require unreliable *sellers*. The proposal, as the I have limited it, does not allow people to act akratically. Buyers have to believe the seller is acting immorally, and the seller has to believe that what they are doing is permissible but not morally good. The moral market

---

[48] Toby Ord 2015 makes a similar point, at 138:

> Indeed if our views on morality are so misguided that we are systematically pointing away from the objective moral requirements, then we would presumably have bigger problems than moral trade, and it would be difficult to blame moral trade for our situation as in that case practical rationality, moral education, and strength of will, would all be making things worse.

does not allow the buyer to pay the seller, or the seller to accept money, to do something that either party believes the seller ought not.

One last response to this line of objection is regulation. An argument for the moral market does not require that every conceivable trade be a good idea; it may require regulation. Even the market for bread, for example, is regulated.[49]

In a series of articles, Glen Weyl and Eric Posner argue that our system of one-person, one-vote is inefficient because it does not allow people to register the intensity of their preferences, and that this allows an indifferent majority to outweigh a passionate minority.[50] To solve this problem, Weyl and Posner argue that society should instead adopt quadratic voting, a system in which people buy as many votes as they want at the cost of the number of votes, squared. Quadratic voting both allows people to register the intensity of their preferences and prevents the wealthy from having too much control: assuming one dollar per vote, it would cost $10,000 for one person to buy 100 votes, whereas it would cost 101 people only $101, in total, to outvote the wealthy voter. Some similar procedure could work for regulating the market for entitlements. How exactly this would work is beyond the scope of this chapter.

Writing on a different, but related topic, Eric Posner and Cass Sunstein argue that when regulatory agencies conduct cost-benefit analyses, they fail to take into account

---

[49] *See,* e.g., 21 Code of Federal Regulations § 136.110 (2005). (Specifying requirements for bread, rolls, and buns.)

[50] *See* Posner and Weyl, "Quadratic Voting and the Public Good," 17; Posner and Weyl, "Voting Squared," 270.

peoples' moral commitments. They argue that people suffer a welfare loss because of the suffering of *others*, and that this welfare loss ought to be taken into account, as well, in regulation.[51] Further, they argue that this welfare loss can be estimated by peoples' willingness to pay to avoid the state of affairs they detest.[52]

This is all to say that regulation may be a powerful tool in responding to problems with the moral market, just as it is for the problems with the market for any good. It is odd, though, especially in a regulatory framework, to think about the welfare loss of one person because of the damage to another. This oddness has been discussed in the literature on the right to destroy: scholars have debated whether the right to destroy persists even if destroying property would cause a welfare loss in others, say, by destroying a work of art,[53] or killing a companion animal.[54] This theme will be taken up in Section 3.4.3, when I discuss an objection to the moral market: that it could lead to extortion.

3.4 Objections

Before beginning, note that the moral market's desirability does not require that every possible trade be wholly unproblematic: some trades may be prohibitively costly, or unwise, or repugnant. Those cases ought to be either regulated or prohibited. Other trades

---

[51] One might expand their thought about the welfare loss people suffer because of the (knowledge of the) pain and suffering of others to include, not only the pain and suffering of others, but knowledge of immoral activity, generally speaking.

[52] Posner and Sunstein, "Moral Commitments in Cost-Benefit Analysis." They note, however, some reasons to be wary of willingness-to-pay analyses, at 9.

[53] *See* Strahilevitz, "The Right to Destroy," 828.

[54] *See* Sachs, "Saving Toby," 251.

may be good on the whole, even if there are costs. As mentioned above, a market for bread, for example, does not require the sale of *all* bread, and does not require that the market be problem-free. This project merely provides the framework for encouraging and regulating moral market trades; it does not justify every possible trade.

3.4.1 THE LIMITS OF THE MARKET

The scholarly literature includes many different arguments for including or excluding certain goods and services from trade. The aim in this section is to briefly explicate the best exclusionary arguments to show that even if one or more arguments for excluding, say, organs, from the market are sound, similar arguments would not work against most trades on the moral market.

One might object that moral decision-making is one of a number of activities that should not be subject to market forces. One might worry that alienating moral deliberation violates an important feature of autonomy.[55] This objection extends familiar arguments against sex work and the sale of human organs, among other exchanges, to moral decision-making. Scholars have articulated different versions of the anti-commodification and coercion objections.[56] Michael Sandel, for instance, argues that market values are

---

[55] *See*, e.g., Kuflik, "The Inalienability of Autonomy." (Discussing the alienability of moral autonomy); Barnett, "Contract Remedies and Inalienable Rights." (Discussing alienable rights and contracts). Scanlon, "A Theory of Freedom of Expression," at 215:

To regard himself as autonomous…a person must see himself as sovereign in deciding what to believe and in weighing competing reasons for action. He must apply to these tasks his own canons of rationality, and must recognize the needs to defend his beliefs and decisions in accordance with these canons.

[56] *See* Anderson, *Value in Ethics and Economics*.

"corrosive" of certain goods. When some goods are subjected to market forces, he writes, "markets change the character of the goods and social practices they govern."[57] He describes how a market-based system of offsets can make some people think that they can absolve themselves from immoral actions, and gives as examples websites for paying for offsets for carbon use and for cheating on exams.[58] In the latter case, people who cheat on exams think they can restore moral balance by making donations.[59]

Margaret Jane Radin argues for the market-inalienability of certain goods and services, while rejecting both universal commodification and universal noncommodification.[60] According to Radin, market-inalienability is justified to protect those things that are important to personhood.[61]

Sandel and Debra Satz both argue that *fairness* should constrain what markets we allow.[62] Satz describes the market for organs, for example, in which sellers are often worse off for the loss of their organs, and for the (sometimes dangerous) procedure to procure them. The sellers are usually poor, and have few choices. Furthermore, markets for organs and the like "undermine the social framework needed for people to interact as equals, as

---

[57] Sandel, *What Money Can't Buy: The Moral Limits of Markets*, 2009, 120.

[58] *See* Foerster, "Moral Offsetting." (Offering an illuminating discussion of moral offsetting.)

[59] Sandel, *What Money Can't Buy: The Moral Limits of Markets*, 2009, 77.

[60] *See generally* Radin, "Market-Inalienability." (Arguing that thinking about human flourishing can provide insight into the limits of the market).

[61] Radin, *Contested Commodities*, 193.

[62] Sandel, *What Money Can't Buy: The Moral Limits of Markets*, 2009, 110–11.

individuals with equal standing."[63] Satz says that this is indicative of *obnoxious markets*.[64] One might, with Sandel and Satz, think that transactions made under coercive circumstances, in particular, undermine fairness and equality.

Subjecting things like sex and organs to market forces may be problematic because the sellers in those cases are often coerced. The thought is that buyers of sex and organs take advantage of the poverty and desperation of sellers in those markets. The extent of coercion is so great, the argument goes, that the would-be sellers are better off if the market for such goods did not exist.

Even if these arguments succeed against the sale of child labor, or organs, or sex, they do not succeed against all transactions on the moral market. It is incorrect to characterize the trade as the sale of moral autonomy. In the vegetarian case, the meat-eater believes that it's neither morally good nor morally bad to continue eating meat. The seller has already made up their mind about the morality of eating meat, and has decided that it is morally neutral.

The heft of any coercion argument is that the seller, but for their dire circumstances, would not engage in the trade because of the harmful or degrading nature of the sale, or because they lack information. This might be right for many instances of sex work or the sale of organs, but in the vegetarianism case, this objection does not work because the seller regards the switch to vegetarianism as, at worst, an inconvenience. That the *buyer* regards

---

[63] Satz, *Why Some Things Should Not Be for Sale*, 95.

[64] Satz, 93–95.

the switch as morally obligatory should not bear on whether the offer of a trade would coerce the buyer—at least, not more than any other service that sellers are morally indifferent toward, like painting a house or preparing a tax return.

On the other hand, cases that *would* involve paying people to do what they believe is morally impermissible *would* be susceptible to these worries. These are Type-2 cases, discussed above. They are excluded from the proposal. Even if there were some price that Bethany would be willing to accept to stop praying, or to pray to a different god, one might think that, because prayer is important to personhood, or human flourishing, or because a market to get people to stop praying would be coercive, one ought not facilitate a market for this type of exchange.

3.4.2 DOING THE RIGHT THINGS FOR THE WRONG REASONS

A Kantian might object that if a seller ceases moral behavior for money, not from a sense of duty, then the cessation lacks moral value. So, participating in the market for entitlements is not morally good. But first, even if the seller's action lacks moral value because they become vegetarian for money and not, say, because they are convinced by Tom Regan,[65] the action of the *buyer* may yet have moral value. The buyer, we might think, pays the seller to go vegetarian out of their sense of duty. Thus, even on a Kantian picture, it seems that trades on the moral market could have moral value.

---

[65] *See generally*, Regan, *The Case for Animal Rights*. (Giving a now-famous rights-based defense of animals rights.)

Even if the Kantian is not persuaded, a trade's lacking moral value is a not a reason *against* it. Kantians presumably think that buying a chocolate bar lacks moral value, and yet Kantians do not generally object to buying chocolate. A Kantian might press that exchanges on the moral market are a violation of the Principle of Humanity.[66] But, as argued above, if one constrains the moral market to just those cases in which the seller regards their behavior as permissible but not morally good, then it is difficult to see how paying someone to become vegetarian constitutes using them as a mere means any more than paying them to paint a house. Were the buyer paying the seller to engage in behavior that the seller themself regarded as immoral, then a plausible case could be made that the exchange violates the Principle of Humanity.

Others may make similar objections. Sandel's warning about the market's corrosive effect on some goods could be a serious problem. As he puts it, the presence of financial incentives may "crowd out" moral or civic norms.[67]

Samuel Bowels is similarly concerned that market values will crowd out moral values. He argues that economists and policy-makers are wrong to think that people are "entirely self-interested and amoral." Rather, he argues, there is a substantial body of psychological studies that show that while people care about acting ethically, when given financial incentives to do so, those financial incentives crowd out the intrinsic moral norms

---

[66] *See generally*, Kant, *The Groundwork of the Metaphysics of Morals*.

[67] *See* Sandel, *What Money Can't Buy: The Moral Limits of Markets*, 2009, 16; Sandel, "How Markets Crowd Out Morals."

people already have.[68] Here, crowding out can mean either replacing the intrinsic ethical motivation many of us have, or it could mean that our moral decision-making would be adversely affected.[69] Among other possible explanations,[70] Bowles describes "moral disengagement," a phenomenon in which "moral reasons become less salient," in particular, when financial incentives are present.[71]

Even if Bowles and Sandel are right that encouraging ethical behavior through financial incentives is problematic, the way the proposal is restricted avoids these concerns. It may very well be problematic to pay people to behave as they might otherwise, out of goodwill, or as they know they ought to. But if the concern is that incentives crowd out morals, the morals have to be there in the first place. On the moral market, buyers pay seller to cease or engage in behavior toward which the seller is morally indifferent. One cannot then say that paying an omnivore to be vegetarian crowds out what otherwise would have been moral motives to become vegetarian, because no such motives exist. Thus, while the crowding out objection may provide a good reason not to pay people to do what they know they ought to do, it does not provide a good reason not to pay people to do what they are morally indifferent toward.

---

[68] Bowles, *The Moral Economy*, 41.

[69] Bowles, 21.

[70] Bowles also mentions dual process theory as an explanation. He says that the presence of incentives stimulates deliberation, which, some scholars think, is bad for moral reasoning.

[71] Bowles, *The Moral Economy*, 96. *See generally* Bandura, "Moral Disengagement in the Perpetration of Inhumanities."; Bandura, "Social Cognitive Theory of Moral Thought and Action." (Explaining the phenomenon of moral disengagement.)

3.4.3 MORAL EXTORTION

One might be concerned that the moral market creates an opportunity for *moral extortion*. Imagine that an omnivore wants to get higher than market price for their entitlement to eat meat, so they artificially inflate its value by threatening to buy hundreds of pounds of meat and throw it away unless someone pays them a great deal not to. Or, imagine that the non-recycler threatens to rent out the local Hertz' fleet and let the cars idle in the parking lot. Or, the non-theist threatens, not only to fail to pray to God, but to pray to the devil. The concern is that by commodifying the entitlement to engage in behavior others believe is morally impermissible, the moral market allows the opportunist to engage in moral extortion. The opportunist preys on the anxiety that their legally permitted behavior will cause others to feel.

One might think that if each individual *owns* the right to engage in legal behavior that others find immoral, one can do anything one wants with the entitlement: after all, the foregoing has argued that readers should think about it as something they can sell. If this were the case, it would be a powerful (though not necessarily decisive) objection against the moral market. Fortunately, there is reason to think that this kind of behavior is, or could be, illegal.

It could be, but is unlikely, that moral extortion is, in some cases, extortion as it is legally understood, according to federal[72] or state law.[73] Under the Hobbes Act, extortion is the "consensual obtaining of property from another induced by wrongful use of actual or

---

[72] Hobbs Act, 18 U.S.C. §1951 (2000).

[73] Model Penal Code § 223.4 (American Law Institute 2017).

threatened force, violence, or fear, or under color of official right." It is often invoked to prosecute public officials for corruption. There is reason to think that moral extortion might not quite fall under the legal meaning of extortion as it is currently understood: the courts would have to interpret "fear" very broadly. Nevertheless, scholars have offered different arguments about how this kind of behavior may yet be illegal. Broadly, there is wide agreement on the following: a person's *reasons* for using their property in a particular way can affect whether they are legally allowed to do so.[74] As Stephen Sachs puts it, "the law may legitimately hold that it is unlawful to threaten maliciously what would, under other circumstances, be entirely lawful to do."[75]

Sachs argues for the creation of a new kind of extortion: extortionate destruction. He describes the odd case of Toby the Bunny: in 2005, someone made a website with photos of his rabbit and threatened to kill and eat him unless people sent him $50,000.[76] While this might appear extortionate, extortion statutes protect people from threats against *their* property. And, while it's possible that the court could construe the dislike that people have of others killing rabbits as "fear," this would, according to Sachs, be overbroad.[77] Indeed, even if we ought to take seriously the welfare costs of Bethany's immoral actions

---

[74] *See* Kelly, "Strategic Spillovers"; Katz, "Spite and Extortion"; Strahilevitz, "The Right to Destroy"; Fennell, "Adjusting Alienability"; Perillo, "Abuse of Rights"; Sachs, "Saving Toby."

[75] Sachs, "Saving Toby," 260. (Adding that this is what happens in the case of blackmail).

[76] Sachs, 251.

[77] Sachs, 251–52. ("Current extortion statues, however, generally do not prohibit the threatened destruction of one's own property, even if they prohibit endangering property owned by someone else. The law thus provides insufficient protection to a variety of resources on which others might place value, including historical buildings, treasured paintings, and adorable bunny rabbits.")

on Adam,[78] it is not the case that any time someone does something others dislike, even for moral reasons, that behavior is or should be illegal. However, when done *specifically* to cause harm, or to get money, what would otherwise have been legal activity, can be illegal.

Moral externality is not altogether a new concept. Daniel Kelly discusses what he calls "strategic spillovers," which are negative externalities intentionally generated by a party's "use of property to extract payments from victims in exchange for desisting."[79] He describes how opportunistic parties generate strategic spillovers in many areas of law: environmental law, intellectual property law, corporate law, legislation and regulation, and litigation and settlement. He does not offer one solution to this problem, but indicates that the most promising might be the abuse of right doctrine.[80]

According to the abuse of right doctrine, while in general people have freedom to use their property as they see fit, when owners use their property precisely to cause harm, either as a means to extract payment or simply for spite, this constitutes an abuse of the right of ownership, and in so doing the owner "exceeds her jurisdiction." As Larissa Katz puts it, "owners lack the jurisdiction to exercise their authority just for the reason that it will cause harm to another."[81]

---

[78] *See* Posner and Sunstein, "Moral Commitments in Cost-Benefit Analysis."

[79] Kelly, "Strategic Spillovers," 1644.

[80] *See* also Kelly, 1710. (Noting that "Unlike the United States, many civil law countries attempt to address explicitly the type of opportunism inherent in strategic spillovers." And continuing, "Under the abuse of right doctrine, a court may prohibit an individual from engaging in what would otherwise be a valid exercise of a legal right if the person is exercising the right for an illegitimate reason.")

[81] Katz, "Spite and Extortion," 1468.

In some cases, the courts have found strategic spillovers to be illegal, if not extortionate. So-called "spite fences" are just one illustration of this. Generally speaking, people can build fences on their property even if their neighbors find them ugly, or block their views, or their sunshine. However, when built for the *purpose* either to spite the neighbor, or to extract payment, building them can be illegal.[82] Katz argues that the court's willingness to find these practices illegal is both explained and justified by the *abuse of right* doctrine. Even if the abuse of right doctrine is not explicit in the common law, Katz and others have argued that the doctrine permeates American law, nevertheless.[83]

Even though the prospect of moral extortion is daunting, there is reason to be hopeful that this isn't too great a problem for the moral market. First, in many cases it may be illegal, as are spite fences and blackmail. Second, even if not illegal now, many have argued that similar cases not currently illegal, ought to be. Were trades on the moral market to flourish, there would be increased pressure on courts and legislatures to protect against moral extortion. Third, that there are problems with the moral market is not decisive against it. Kelly describes the existence of strategic spillovers in many areas of the law, and yet the solution to the problem is not to eliminate all of the very many kinds of interactions that bring it about.

---

[82] *See* Kelly, "Strategic Spillovers," 1667–68; Sachs, "Saving Toby," 259–60; Fennell, "Adjusting Alienability," 1454–55.

[83] *See* Perillo, "Abuse of Rights."

3.4.4 UNCONSCIONABILITY

The proposal excludes both paying someone to engage in and desist from illegal activity. The proposal has assumed that contracts on the moral market would be legally binding and enforceable. However, some contracts, though legally valid, will not be enforced in courts of equity if they are *unconscionable*.[84] While not precisely defined, a contract is said to be *unconscionable* if "it was 'such as no man in his senses and not under deluston would make on the one hand, and as no honest and fair man would accept on the other[.]'"[85] The court employs a two-pronged test to determine whether a contract provision is unconscionable—it tests for both procedural and substantive unconscionability.[86] The former tests for unfairness in the way the contract came about; the latter tests for unfairness in the exchange itself. Unconscionability has been invoked in instances of vastly unequal bargaining power,[87] where all of the parties offering the service include the same clause,[88] and where it provides for under-compensation in the case of a

---

[84] *See* Shiffrin, "Paternalism, Unconscionability Doctrine, and Accommodation," 207. (Defending the doctrine against a charge that it is overly paternalist); *see also* Bagchi, "Distributive Justice and Contract," 193. ("[Principles of distributive justice] are among the moral considerations that appropriately inform rules of validity, interpretation, and remedy." He continues, "[c]ontracts that take place against a backdrop of distributive *injustice* may be subject to further…constraints.")

[85] Restatement (Second) of Contracts § 208 cmt. b (American Law Institute 1981) (citing *Hume v. United States*, 132 U.S. 406 (1889)).

[86] *See* McCullough, "Unconscionability as a Coherent Legal Concept," 781. (Providing an excellent overview of the topic.)

[87] U.C.C. § 2-302 cmt. 1 (American Law Institute 2018) ("The principle is one of the prevention of oppression and unfair surprise . . . . ")

[88] *Williston on Contracts*, sec. 18:13.

breach.[89] It has also been invoked where one of the contracting parties doesn't know what they're doing.[90] In general, contracts have been found to be unconscionable where they are grossly unfair.[91]

Some trades on the moral market might be susceptible to conscionability worries, but the likelihood is low, both because the standard for unconscionability is high, and because the market is limited to those trades in which the seller believes they are being paid to do something toward which they are morally indifferent. Where the contract *is* abhorrent to the court, it may, for unconscionability reasons, refuse to enforce the terms. Given the constraints enumerated, there is reason to think that these cases will be rare.

In addition, the unconscionability doctrine may be a tool for regulating the unsavory borders of the moral market. Some argue that the doctrine reflects society's unwillingness to endorse socially destructive agreements, especially when those agreements may undermine equality.[92] Debra Satz makes a related point, arguing that the threat to equality is one indication of an "obnoxious market."[93] And, it is well-documented both that the

---

[89] *See, e.g.*, Fried, "The Ambitions of Contract as Promise," 27. (Citing U.C.C. § 2-718 cmt. 1. (American Law Institute 2018)).

[90] *Williston on Contracts*, sec. 18:8; Radin, *Contested Commodities*, 228.

[91] *See* Cornell, "A Complainant-Oriented Approach to Unconscionability and Contract Law," 1146; Shiffrin, "Paternalism, Unconscionability Doctrine, and Accommodation," 209.

[92] *See, e.g.*, Bagchi, "Distributive Injustice and Private Law," 109; Shiffrin, "The Divergence of Contract and Promise," 752.

[93] Satz, *Why Some Things Should Not Be for Sale*, 95. (Explaining that "[t]he operation of these markets can undermine the social framework needed for people to interact *as equals*, as individuals with equal standing"); *see also,* Quong, "Cultural Exemptions, Expensive Tastes, and Equal Opportunities," 57. ("I believe that principles of justice should have a fundamentally different purpose, which is to create a system of regulation for political society that enables all persons to live as free and equal citizens.").

courts are increasingly willing to refuse to enforce contracts as unconscionable, and that more litigants are bringing claims.[94] Parties are less likely to form contracts they doubt courts will enforce.

3.4.5 THE WEALTH EFFECT

The market allocates resources more efficiently than their initial allocation. When Adam values Bethany's entitlement to eat meat more than she does, a trade leads to greater efficiency (absent market failure). And, intensity of preference determines the allocation of goods: the person who would derive the most from a good ends up with it, at least in theory. But, one might worry that, as Saul Levmore puts it, "wealth effects might dominate preference intensities."[95] Call this the *wealth effect*.[96]

Of course, the wealth effect is not limited to the moral market. It is problematic, also, in the market for concert tickets, health care, and diamond jewelry. Is it then, *especially* problematic in a way that constitutes an objection to the proposal? The proposal is already limited to cases where the buyer pays the seller to cease behavior toward which the seller is morally indifferent. Thus, one need not worry that the wealth effect would induce a poor seller to do something they believe they ought not. Perhaps, then, the worry

---

[94] *See* McCullough, "Unconscionability as a Coherent Legal Concept," 786–87.

[95] Levmore, "Voting with Intensity," 160.

[96] *See* Sunstein, "Incommensurability and Valuation in Law," 849; Dagan and Fisher, "Rights for Sale," 98–99; Levmore, "Voting with Intensity," 118. ("Where wealth differentials are present, A might buy B's vote (if legally permitted to do so) even though B is an equal or higher valuing user of that vote. And wealthy people like A might systematically favor different political outcomes than would people with endowments more like B's, so that there is at least an argument for barring trades despite the fact that the buyer and seller are made better off than before. This argument applies to some but not most other inalienable commodities.")

is that the wealth effect will crowd out poor *buyers*. That this is unfair is neither unique to the moral market nor a good objection to it. Poor buyers are crowded out from many markets, after all.

Perhaps the worry is that the buyers will be comparably wealthy, and that the wealthy, on average, have impoverished moral beliefs. It is unclear why we should think this is the case. Any explanation that makes reference to the wealthy's poor behavior might be better explained by enhanced opportunity to act akratically.

The best objection invoking the wealth effect is that the wealthy would have comparably greater control of the moral landscape, which is inegalitarian even *if* they are no less likely to have the correct moral views.[97] Note that this seems to be what happens already.[98] And, while many charitable donations go toward aid, such as food, medicine, disaster relief, and the like, this is not the case for all charities. Consider any charitable organization that, in part or in whole, is devoted to changing people's minds. If the wealth effect is problematic for the moral market, it seems similarly problematic for at least some kinds of tax-deductible charitable giving.

However, the proposal for a moral market is not a proposal for an *unregulated* market. The market ought to be limited—to exclude, for example, paying people to do things they believe they ought not, and to exclude those cases likely to involve coercion.

---

[97] *See* Stilz, "Is the Free Market Fair?," 432.

[98] *See generally,* Mayer, *Dark Money*. (Describing the machinery of lobbying by America's wealthiest residents).

As Saul Levmore writes, "…there are ways of limiting wealth effects, and clever ideas for limited markets may soon begin to surface."[99]

As discussed in Section 3.3.4, quadratic voting could work as a check on the influence of the wealthy, while at the same time allowing people to register the intensity of their preferences. As will be discussed in Section 3.5.2, market design may prevent some of the wealth effect: markets for kidneys, for example, operate without money, and we are all, roughly speaking, equally endowed with kidneys. As Tsilly Dagan and Talia Fisher emphasize, imposing different kinds of inalienability mechanisms can promote efficiency while blunting the effects of inequality. They suggest a modified inalienability rule, unrestricted barter, for school vouchers, for example.[100] Further methods of regulation include progressive taxation on trades on the moral market, or transfer payments.

3.4.6 TRANSACTION COSTS

There is reason to wonder whether this proposal is feasible, given transaction costs. Finding willing buyers and sellers and drawing up and enforcing contracts might make these trades too expensive. If a buyer pays a meat-eater to stop eating meat, how will they know whether the seller honored the contract? Call this the *enforcement problem*. That this problem exists, however, is not decisive against the moral market. That *some* trades will be too costly does not undermine the proposal altogether. Transaction costs make lots of trades that would otherwise be wise, unwise.

---

[99] Levmore, "Voting with Intensity," 160.

[100] Dagan and Fisher, "Rights for Sale," 114–16.

For some trades the enforcement problem would be small. Imagine a trade where the buyer pays the seller to go to church. Enforcement would be a matter of taking attendance, or checking in on one's phone using GPS tracking. Generally speaking, then, trades on the moral market ought to be limited to those cases in which the transaction costs are not prohibitive. I discuss ways to decrease transaction costs in Section 3.5.1.

3.5 Two Similar Proposals

If the moral market is a good idea, one would have to come up with a plan. I'm inclined to punt to entrepreneurs, choice architects, and the like.[101] However, I anticipate that many will dislike the proposal, perhaps not only because of objections to specific parts of the argument, but rather because of a general dislike toward the nature of the proposal: paying people to behave how they ought to. The remainder of the article, then, will consist of two, similar proposals, in the same spirit, but which may be more palatable for those put off by the idea of selling entitlements.

3.5.1 PAY TO PRAY

Instead of paying people to cease immoral behavior, one might instead pay people to do the thing that would *convince* them to cease immoral behavior. For example, the vegetarian, instead of paying the omnivore to give up meat, could pay the omnivore to read *Animal Liberation*,[102] or watch a documentary about animal agriculture, in the belief and hope that some people would thereby become convinced to become vegetarian. Note that

---

[101] *See generally* Thaler and Sunstein, *Nudge: Improving Decisions about Health, Wealth and Happiness*. (Arguing that because people behave in predictably irrational ways, those with power can change defaults and thereby make better choices easier to make).

[102] Singer, *Animal Liberation*.

this tactic is employed by VegFund, which, provides money to animal rights groups to pay people to watch videos of the workings of animal agriculture.[103] Mercy for Animals and Compassion for Animals have also employed this tactic.[104] Arguably, this is also part of the motivation for Birthright Israel, an organization that takes young Jews on a free trip to Israel to "…motivate young people to continue to explore their Jewish identity and support for Israel…."[105]

This method would be a valuable tool especially in cases where it may be impossible to pay someone to do what we believe is morally obligatory. For example, some believe that it is morally obligatory to convert to Christianity. However, these same people might believe that conversion requires, say, belief that Christ is the lord and savior, and that, because belief is involuntary, you can't simply pay people to be Christian. But, one could pay someone else to do the things that would maximize that person's chance of acquiring the necessary beliefs, and then, converting: say, attending church services, or reading the Bible.

If someone who hands out religious literature thinks that someone's reading it will increase their chance of converting, then increasing the percent of people who read it would presumably increase the number of converts. For example, outside my office, from time to time a member of a local church hands out granola bars wrapped in paper on which Bible

---

103 VEGFUND, https://vegfund.org/category/grant-programs/paid-per-view-ppv/

104 Nathan Runkle, Mercy For Animals, "Farm to Fridge Inspires Students to Boycott Factory Farm Cruelty."

105 "*Our Story,*", Birthright Israel. [https://perma.cc/6MGE-USUK].

verses and the time and location of church services are written. Each time a student takes a granola bar, they are being paid with a snack to have to at least glance at a Bible verse and information about the church service. This church member must believe that people will be more inclined to read the note if it is wrapped in a granola bar. On the same quad, a student group offers an electronic coupon for a free coffee to any student who sends, as a text message to the group, a "question you have for God." The student thereby joins the mailing list, and, perhaps, gets an answer to their question, presumably from a student in the group. Here, the group is paying students with drinks to be exposed to their emails, in the hope that some students will become interested, and join. And perhaps, convert.

The weakness of this alternative proposal is that it does not guarantee compliance. On the moral market, compliance is contractual. However, this proposal does have several benefits. First, it is a good alternative when it is not possible to pay someone to engage in the desired behavior, itself, as in the religion case, above. Second, it more closely resembles the deliberative process some might feel should uniquely determine the means by which we make moral decisions.[106]

Not only this, but the buyer might have reason to believe the change will be longer-lasting. If the buyer pays someone to become a vegetarian for a week, the seller has no reason to continue after the contract expires. However, if the seller is paid to read

---

[106] But *see generally* Haidt, *The Righteous Mind: Why Good People Are Divided by Politics and Religion.* (Arguing that there is reason to doubt that our moral beliefs are formed exclusively through deliberative processes.)

arguments in favor of vegetarianism, there is a greater likelihood that some percent of sellers will become vegetarians for longer than they would have, otherwise.

Third, the transactions costs would likely be lower: imagine the difference between tracking compliance on whether someone remains a vegetarian for a week, on the one hand, and whether they have read a book or watched a film, on the other. On a website, it would be easy to provide an article or film about vegetarianism, or the environment, or Christianity, and then administer a quiz to test understanding. Certainly Silicon Valley could create such a platform. Or, one could use existing infrastructure and methodology. Human Resources and IT offices routinely provide training through videos about various company policies: harassment, FERPA compliance, copyright, etc. Amazon Mechanical Turk, for example, is routinely used to find and pay people willing to participate in all sorts of tasks.

For a forum in which to present the information, Reddit's "Change My View" might provide a model.[107] On this website, someone poses a question, and then various commentators give arguments for or against. People then indicate if they change their minds. It has roughly half a million subscribers.[108] This shows that some people are willing, for free, to be convinced that they are mistaken.

Pay to Pray more closely resembles the way we normally try to convince people to change their views. It is difficult, however, to get people to listen; the attention of the public is a valuable commodity, after all. By compensating people for their time, they would

---

[107] "Change My View (CMV) • r/Changemyview."

[108] "*/r/changemyview metrics*," [https://perma.cc/7T4F-3VBL].

theoretically be more willing to engage in a deliberative practice similar to the one that the buyer engaged in, initially.

3.5.2 MARKET DESIGN

For those opposed to paying others to engage in or desist from behavior, for moral reasons, alternative market design might yield an answer. Economists have written extensively on *repugnance* and the way that it has and continues to constrain markets. Repugnance is meant to reflect, as Julio Jorge Elias calls it, the "yuck factor" and is distinct from a belief that the good or service traded would cause harm, as it might with some drugs, or guns.[109] Alan Roth offers repugnance as the explanation for California's ban on horse meat, for example.[110] Thus, an argument that a market in some good would make people better off is not sufficient to overcome the repugnancy objection. In some cases, the transactions are repugnant precisely *because of* the introduction of money. Roth writes that, although transactions can be repugnant for different reasons, "[o]ne often-noted regularity is that some transactions that are not repugnant as gifts and in-kind exchanges become repugnant when money is added."[111] Sometimes, then, the repugnance can be overcome through strategic market design, for example, by replacing financial transactions with in-kind payments.[112] This has happened in the market for kidneys.

---

[109] Elias, "The Role of Repugnance in the Development of Markets," 234.

[110] Roth, "Repugnance as a Constraint on Markets," 37–38.

[111] Roth, 44.

[112] Roth, 52. (Noting that while an article suggesting modest payments for organ donation received a negative reaction, his proposal for kidney exchange did not).

Thousands of Americans die every year while waiting for kidneys.[113] While under some conditions people can donate organs, their sale is prohibited in the United States.[114] Were there a legal market for kidneys, fewer people would die for want of a transplant. The lack of a (legal) market is problematic not only for people who have no willing donors, but also for those in need of an organ with willing but medically incompatible donors. Economists designed a market without money to solve the problem of willing but incompatible donors, without changing the law.

"Kidney exchanges" allow a workaround to the repugnancy problem. In a one kind of exchange, a *paired donation*, Patients A and B each have willing but medically incompatible donors. Through a database, doctors determine that A's donor is compatible with B, and B's donor with A's, and carry out the surgeries accordingly. Thus, it is as if A's donor has given to A, when in fact their kidney went to B. This solves two problems: first, it allows people who otherwise could not to donate organs to their family members; second, it does so without using money.[115] In another kind of exchange, a non-directed donor gives a kidney and the reciprocity comes at a future, uncertain time, from an uncertain source—when a compatible kidney becomes available, their friend or relative is

---

[113] UCSF, "The Kidney Project." In 2014, over 100,000 patients were on the transplant waitlist, and only roughly 17,000 donor kidneys were available.

[114] National Organ Transplant Act, 98 Stat. 2339 (1984). Using Rose-Ackerman's terminology, organs are controlled with a *modified* inalienability rule because they can be given, but not sold. In contrast, votes are controlled with a *pure* inalienability rule, because one can neither sell nor give away their vote (at least in elections for government positions).

[115] *See* Roth, "What Have We Learned from Market Design?," sec. III; Roth, "Repugnance as a Constraint on Markets," 45–52.

entitled to it.[116] The chain could involve a large number of donations. Alvin Roth explains it this way:

> …exchanges could be a cycle of incompatible patient-donor pairs of any size such that the donor in the first pair donated a kidney to the patient in the second, the second pair donated to third, and so on until the cycle closed, with the last pair donating to the first.[117]

Kidney exchange might provide a model for getting around repugnancy worries about the moral market. Rather than buying and selling entitlements, people could trade them through either paired or chain "donations." If the vegetarian cannot *buy* the omnivore's entitlement to eat meat, they could, instead, trade, either through a paired donation, if the two parties are compatible, or through a chain donation, if they are not. There might be something that the vegetarian does that the omnivore believes is morally impermissible and toward which the omnivore is morally indifferent.[118]

This is all meant just by way of example, and not as advocacy for a particular kind of market design. Trades are a *kind* of payment, of course, but they get around some repugnancy problems, and might be a good response to an objection that one might have thought inherent to the moral market.[119]

---

[116] *See, e.g.*, Woodle et al., "Ethical Considerations for Participation of Nondirected Living Donors in Kidney Exchange Programs," 1461.

[117] Roth, "Repugnance as a Constraint on Markets," 86.

[118] Choi, Gulati, and Posner, "Altruism Exchanges and the Kidney Shortage," 292. (Making a similar proposal to alleviate the kidney shortage by proposing trading kidneys for non-kidney altruistic donations.)

[119] *See, e.g.*, Hasen, "Vote Buying," 1339. (Describing this feature of Congressional logrolling).

3.6 Conclusion

This paper presents what some may consider a radical proposal: that people ought to pay others to do what they themselves believe others ought to do. In some ways, though, the view is not all that radical—people pay others to do what they want them to do all the time. Properly limited, and regulated, the moral market would allow parties to trade in a way that makes them better off, and, according to at least one party, things would be morally better. In addition, one has reason to believe the market would contribute to actual moral progress. The intuition that moral decision-making ought not be commodified is, in this author's view, the right one. However, the kinds of trades permitted by the proposal cannot properly be described as paying someone to give up their moral autonomy. Paying someone to do something they are morally indifferent toward is as common a feature of the market as the sale of bread.

As said at the outset, the proposal is highly speculative. The purpose of this paper is to draw attention to the inefficiency of one kind of entitlement, and to suggest a framework for addressing it. That this proposal is imperfect, or that the moral market would have problems, is not definitive against it.

4 AGAINST THE ALLEGED INSUFFICIENCY OF STATISTICAL EVIDENCE IN TRIALS

4.1 Introduction

We should take into account the relevant evidence. This is a general principle in epistemology and in evidence law scholarship, and is expressly stated in the Federal Rules of Evidence.[120] Where the evidence is, or is likely to be, inaccurate, it should be discounted or excluded from consideration. And the same where the evidence is accurate but likely to mislead.[121] In some cases, there is controversy about what kinds of evidence are likely to be inaccurate, or likely to mislead. But, it is uncontroversial that *when* evidence is inaccurate or misleading, this provides a strong if not decisive reason for excluding or discounting that evidence: reasoning with such evidence inhibits factfinding. Call these reasons for exclusion *accuracy-based* epistemic reasons. At other times, the weight we accord evidence is not determined solely by accuracy-based reasons, but for what might be referred to as *policy* reasons. For example, some philosophers have argued that we ought to defer, against the evidence, to friends.[122] In a similar vein, stereotyping may be morally objectionable, even if it is not (always) evidentially deficient.

---

[120] *See,* e.g., Pundik, "The Epistemology of Statistical Evidence," 137.; "Federal Rules of Evidence" (FRE) Rule 402 (Allowing all relevant evidence unless expressly proscribed).

[121] For example, FRE Rule 403 allows the exclusion of relevant evidence "…if its probative value is substantially outweighed by a danger of one or more of the following: unfair prejudice, confusing the issues, misleading the jury…"

[122] *See,* e.g., Stroud, "Epistemic Partiality in Friendship"; Keller, "Friendship and Belief." (Arguing that we should sometimes defer, against the evidence, to friends). But see Kawall, "Friendship and Epistemic Norms." (Responding to Stroud and Keller).

In the literature on the use of statistical evidence in trials, the consensus view among legal scholars and philosophers is that statistical evidence should not be, on its own, sufficient for conviction in a criminal case or a judgment in a civil case. With few notable exceptions,[123] neither legal scholars nor philosophers give accuracy-based epistemic reasons for excluding or discounting statistical evidence. For the most part, philosophers have given *non*-accuracy-based epistemic reasons. That is, statistical evidence is deficient in some epistemic virtue, but this deficiency is not one that would inhibit fact-finding. And, for the most part, legal scholars have given policy-based reasons.

I argue that we should not discount statistical evidence, as scholars have claimed and as the courts have sometimes held. I argue that statistical evidence *should be* sufficient, on its own, for conviction in a criminal trial or a judgment in a civil trial.[124] I start with a presumption that all relevant evidence should receive its due. The main claim of this paper is that the considerations scholars give in favor of statistical evidence's insufficiency do not overcome this presumption. Even if the considerations on offer are compelling, the cost

---

[123] *See,* e.g., Di Bello, "Trial by Statistics: Is a High Probability of Guilt Enough to Convict?"

[124] I'm not the first person to say that non-accuracy based epistemic reasons don't lead neatly to a conclusion that statistical evidence should not be sufficient for a finding of legal responsibility. See, especially, Enoch, Spectre, and Fisher, "Statistical Evidence, Sensitivity, and the Legal Value of Knowledge"; Enoch and Fisher, "Sense and 'Sensitivity': Epistemic and Instrumental Approaches to Statistical Evidence." (Giving policy-based reasons for thinking that statistical evidence's insensitivity provides a perverse incentive to wrongdoers in some cases, and arguing, generally, that the courts should be more concerned with factfinding than with knowledge); Pundik, "The Epistemology of Statistical Evidence," 122, fn 27: "Proponents of the distinction [between admissible and inadmissible statistical evidence] need to provide a more refined distinction between acceptable and problematic statistical evidence, together with some explanation as to why any differential treatment of objectionable statistical evidence should not apply to types of statistical evidence they consider acceptable."; Pundik, "What Is Wrong with Statistical Evidence? The Attempts to Establish an Epistemic Deficiency," 463. (Arguing that none of the epistemic reasons "successfully establish[] an epistemic deficiency from which (only) statistical evidence suffers.")

of demoting statistical evidence is substantial. Broad pronouncements about epistemic

desiderata, I think, draw attention away from the purpose of the trial, which is, at bottom,

a factfinding mission.[125] Of course, this is not the only goal—if it were, then the

Exclusionary Rule, which renders probative evidence inadmissible if it was obtained

unlawfully, would not exist. But often the discussions seem to miss the cost to accuracy,

which in turn is a cost to victims of crimes, their families, to future victims, and to those

who suffer torts. In an article on this topic, Duncan Pritchard makes a claim about error in

trials: "In short, we want a criminal justice system that excludes high levels of risk of

wrongful conviction, where risk is understood modally rather than probabilistically."[126] It

is unclear why anyone affected by the risk of wrongful conviction would prefer a modal

conception of risk, even if it has theoretical virtues. One would think that the people

involved with and affected by trials would care primarily about accuracy. The Court has

alluded to this many times.[127] And, in the absence of sufficient countervailing reasons to

prize non-accuracy-based epistemic virtues over accuracy, it's difficult to see why those

reasons ought to prevail.

---

[125] "The basic purpose of a trial is the determination of the truth." (*Tehan v. U.S.* 383 U.S. 406, 416)(1966))
Quoted in Laudan, *Truth, Error, and Criminal Law: An Essay in Legal Epistemology*, 2.

[126] Pritchard, "Risk," 454.

[127] "Any claim for the exclusion of evidence logically relevant in criminal prosecutions is heavily
handicapped. It must be justified by an over-riding public policy expressed in the Constitution or
the law of the land." *Nardone v. U.S.*, 308 US 388 (1939) at 340; "…this admittedly drastic and
socially costly course is needed to deter police from violations of constitutional and statutory
protections." *Nix v. Williams,* 467 U.S. 431 (1984) at 442; "The exclusionary rule generates
'substantial social costs[.]" *Hudson v Michigan*, 547 US 586 (2006) at 591.

The paper is organized as follows. First, I describe the problem statistical evidence presents in the law. In brief, the following sort of dilemma arises: Factfinders (the judge or jury) in criminal and civil trials are charged with reaching a verdict if the evidence presented meets a particular standard of proof—beyond a reasonable doubt, in criminal cases, and preponderance of the evidence, in civil trials. It seems that purely statistical evidence can suffice for just such a level of certainty in a variety of cases where our intuition is that it would nonetheless be wrong to convict the defendant, or find in favor of the plaintiff, on merely statistical evidence. So, we either have to convict with statistical evidence, in spite of an intuition that this is unsettling, or else explain what (dispositive) deficiency statistical evidence has.

Second, I discuss the reasons philosophers give for statistical evidence's deficiency—mostly non-accuracy-based epistemic reasons. In general, the arguments have the following form: For evidence to be sufficient for knowledge or belief, it has to have one or another epistemic property. Statistical evidence lacks this property for one reason or another. The factfinder must have the relevant belief, or knowledge, in order to convict a defendant or award a judgment to a plaintiff. Therefore, statistical evidence cannot be sufficient for conviction or a judgment. I don't argue that the relevant philosophers are mistaken about statistical evidence's lacking the properties they identify. But I argue that these philosophers focus too much on the first two premises and not enough on the third— that the right kind of belief or knowledge is necessary in the legal context. In perhaps the most egregious instance of this, Sarah Moss argues that criminal defendants have a right that the jury *know* that they are guilty and that statistical evidence, itself, is not sufficient

78

for that kind of knowledge.[128] Even if we grant the second conjunct, which is part of the argument of her sophisticated and well-received book, there is no good reason to accept the first. And Moss offers no such compelling reason.

In the evidence law literature, scholars mostly give what I've called "policy" reasons against the sufficiency of statistical evidence for conviction or judgment: it undermines a defendant's right to be treated as an individual;[129] it can be difficult to determine *which* reference class someone ought to be considered a member of;[130] it requires an overt admission of error;[131] and, for legitimacy reasons, the public has to be able to treat the verdict as a conclusion about something that *happened*, which, we have reason to think, may not be the case if statistical evidence is sufficient for conviction.[132] I won't discuss these at length in the paper—my focus is on non-accuracy-based epistemic reasons. In large part, some of the policy reasons require a kind of empirical analysis that is, for me, either

---

[128] "According to my account, defendants have the right to be convicted on the basis of nothing less than knowledge." Moss, *Probabilistic Knowledge*, 215.

[129] Wasserman, "The Morality of Statistical Proof and the Risk of Mistaken Liability," 943. (But noting that where the statistical evidence does not "involve an inference to the defendant's conduct from the frequency of similar conduct," as in, e.g., fingerprint analysis, to that extent the statistical evidence may be unproblematic.)

[130] Colyvan, Regan, and Ferson, "Is It a Crime to Belong to a Reference Class?"

[131] "There is something intrinsically immoral about condemning a man as a criminal while telling oneself, 'I believe there is a chance of one in twenty that this defendant is innocent, but a 1/20 risk of sacrificing him erroneously is one I am willing to run in the interest of the public's—and my own— safety." Tribe, "Trial by Mathematics: Precision and Ritual in the Legal Process," 1372, quoted in Koehler, "When Do Courts Think Base Rate Statistics Are Relevant?," 337.

[132] Nesson, "The Evidence or the Event?"

too speculative without data or beyond the scope of this project, otherwise. Moreover, it involves a sort of weighing of costs and benefits that I'm not prepared to undertake.[133]

Third, I discuss the way in which the philosophical literature is out of touch with the legal reality—that outside of toy cases discussed (though some are based on actual cases), the courts allow statistical evidence to be sufficient in a variety of contexts. For example, in the civil context, statistical evidence can be sufficient where either the harm is collective or else it is impossible to tell who, exactly, is responsible for the harm. Slightly more controversially, statistical evidence is used, decisively, in bail and sentencing determinations. Of course, that the practice exists in other contexts does not itself justify the practice, but to the extent that these practices have become accepted for principled reasons, it does. The various considerations I advance lead to the conclusion that when it comes to naked statistical evidence, philosophers who argue for its insufficiency are grasping at straws.

4.2 The Problem Statistical Evidence Poses in Trials

4.2.1 THE PARADOX

As a rough approximation, "statistical evidence" in the legal epistemology literatures refers to evidence from which the fact-finder draws a statistical inference to the defendant's guilt or liability. While it may be specious to distinguish "statistical evidence"

---

[133] As Koehler and Shaviro write, the determination "of whether, on balance, greater use of overtly probabilistic evidence and methods at trial is desirable…depends on the value attached to specific policy concerns other than verdict accuracy." Koehler and Shaviro, "Veridical Verdicts: Increasing Verdict Accuracy Through the Use of Overtly Probabilistic Evidence and Methods," 248.

from other kinds of evidence, I'll stick to this terminology for consistency's sake. Scholars in this literature distinguish between *individualized*, or *particularized*, evidence, on the one hand, which is said to be "about" the defendant, and *statistical evidence*, on the other, which is not. Base rates, for example, are described as "statistical evidence." More philosophically sophisticated scholars working in this area have pointed out this suspect classification of different kinds of evidence.[134]

The statistical evidence is usually about a reference class of which the individual is a member. For example, assume that we have no knowledge of Simon's shoe preferences. But he is a late-twenties male, and we have robust data about the shoe-buying practices of late-twenties males. Thus, we have statistical evidence of Simon's shoe preference, even though we don't know anything about Simon, individually. So, without more information, we can infer from the data something about Simon's shoe preferences. In contrast, if Simon had made a statement about his love of Nikes, then we would have *direct* evidence.

In the legal context, a conviction or a finding is warranted when the government or plaintiff meets the relevant standard of proof. And it seems that statistical evidence, sometimes, is sufficient to meet this burden.[135] Why then, does it seem inappropriate to

---

[134] *See,* e.g., Pardo, "The Paradoxes of Legal Proof: A Critical Guide," 262, fn 128; Shaviro, "Statistical-Probability Evidence and the Appearance of Justice," 530 (Challenging this distinction and arguing that it is ill-defined in the literature). Ron Allen makes a more strident comment: "the assumption that there are two qualitatively distinct types of evidence, statistical and non-statistical, is essentially false." Allen, "On the Significance of Batting Averages and Strikeout Totals: A Clarification of the 'Naked Statistical Evidence' Debate, the Meaning of 'Evidence,' and the Requirement of Proof Beyond a Reasonable Doubt," 1093. Judge Posner in *Riordan v. Kempiners*, 831 F.2d 358, 360 (7th Cir. 1998): "All evidence is probabilistic—statistical evidence merely explicitly so." (cited in Koehler 2002, 401, fn 165).

[135] See Pardo, "The Paradoxes of Legal Proof: A Critical Guide," 253.:

convict or find as the evidence suggests, when the evidence is statistical in nature? One scholar writes that doing so would be "dubious."[136] In the relevant literature, a few cases have come to be canonical representations of the problem. I'll give two civil cases and one criminal case, to illustrate :

Blue Bus

> Suppose it is late at night…and an individual's car is hit by a bus. This individual cannot identify the bus, but she can establish that it is a blue bus, and she can prove as well that 80 percent of the blue buses in the city are operated by the Blue Bus Company, that 20 percent are operated by the Red Bus Company, and that there are no buses in the vicinity except those operated by one of those two companies….In these circumstances can the plaintiff recover in civil litigation against the Blue Bus Company, or, if not…then why not?[137]

Gatecrasher

> Consider a case in which it is common ground that 499 people paid for admission to a rodeo, and that 1,000 are counted on the seats, of whom A is one. Suppose no tickets were issued and there can be no testimony as to whether A paid for admission or climbed over the fence. So there is a .501 probability, on the admitted facts, that he did not pay. The conventionally accepted theory of probability would apparently imply that in such circumstances the rodeo organizers are entitled to judgment against A for the admission money, since the balance of the probability would lie in their favor. But is seems manifestly unjust that A should lose when there is an agreed probability of as high as .499 that he in fact paid for admission.[138]

What makes the examples 'paradoxical' is that the evidence appears on its face to surpass the applicable standard of proof and, yet, the judgment of most people is that the evidence is insufficient to prove liability or guilt. This apparent inconsistency between what the applicable legal rules appear to require, on the one hand, and judgments about what the correct result ought to be, on the other, creates a tension that calls out for explanation.

[136] See Gardiner, "Legal Burdens of Proof and Statistical Evidence," 3.

[137] Schauer, *Profiles, Probabilities, and Stereotypes*, 82. (This case is modeled on *Smith v. Rapid Transit, Inc.*, 58 N.E.2d 754 (Mass. 1945)).

[138] Kaye, "The Paradox of the Gatecrasher and Other Stories," 101. (This case was originally presented in Cohen, *The Probable and the Provable*.).

Prison Yard

> In an enclosed yard are twenty-five identically dressed prisoners and a prison guard. The sole witness is too far away to distinguish individual features. He sees the guard, recognizable by his uniform, trip and fall, apparently knocking himself out. The prisoners huddle and argue. One breaks away from the others and goes to a shed in the corner of the yard to hide. The other twenty-four set upon the fallen guard and kill him. After the killing, the hidden prisoner emerges from the shed and mixes with the other prisoners. When the authorities later enter the yard, they find the dead guard and the twenty-five prisoners. Given these facts, twenty-four of the twenty-five are guilty of murder.[139]

In the above cases, the relevant standard of proof—preponderance of the evidence, generally thought to mean more likely than not, in the civil case; beyond a reasonable doubt, in the criminal case—seem to be met.[140] And yet, the intuition of most scholars, and, in some cases, the courts, is that it has not been. Describing the Blue Bus case, for example, Sean Sullivan writes, "Everyone agrees that Bayesian logic compels this result, but no one seriously thinks it is the right outcome."[141] Describing the actual case on which Blue Bus is based, Andrea Roth writes that the court ruled in favor of the defendant because "…a rational prediction by the jury based on the evidence 'was not enough' absent an 'actual belief' in liability 'in the mind or minds of the tribunal.'"[142] Others have described basing

---

[139] Nesson, "Reasonable Doubt and Permissive Inferences," 1192–93.

[140] There is significant controversy about what standard of proofs are meant to be. I largely gloss over these issues.

[141] Sullivan, "A Likelihood Story," 45.

[142] Roth, "Safety in Numbers? Deciding When DNA Alone Is Enough to Convict," 1164. (Quoting *Smith* at 755).

a judgment on purely statistical evidence as "patently absurd."[143] In addition, this intuition is shared by the participants in several psychology studies, which find, in general, that people are reluctant to make liability decisions when the evidence is based on naked statistics.[144]

On the other hand, in what I think is a deficiency of the literature in this regard, there is an undue focus on the defendant—criminal and civil. When scholars (mostly legal scholars) give policy-based reasons for the insufficiency of statistical evidence, they often invoke notions of fairness.[145] That is, fairness to the defendant. Neglected are the costs imposed on the plaintiffs (in civil cases) and on society, and perhaps the victims, (in criminal cases). In the criminal context, at least part of the justification for the criminal justice system is to keep people safe. In a related literature, scholars have debated not only what standards of proof mean, but, more provocatively, whether we are using the right standards. Using estimates of the likelihood of being the victim of a violent crime and being falsely convicted of such a crime, respectively, and the relative harms of each, Larry Laudan argues that we ought to lower the standard of proof in criminal trials.[146] Even those

---

[143] Allen and Stein, "Evidence, Probability, and the Burden of Proof," 574. (Cited in Sullivan, "A Likelihood Story," 13.)

[144] Wells, "Naked Statistical Evidence of Liability: Is Subjective Probability Enough?"; Wright et al., "Factors Affecting the Use of Naked Statistical Evidence of Liability."

[145] Enoch, Spectre, and Fisher, "Statistical Evidence, Sensitivity, and the Legal Value of Knowledge" is a notable exception.

[146] *See,* generally, Laudan, *The Law's Flaws*. For a pointed criticism of Laudan's argument, see Gardiner, "In Defence of Reasonable Doubt."

who do not agree with Laudan's analysis recognize the substantial costs to ignoring evidence, as discussed, above.

In American law, for example, the Exclusionary Rule precludes evidence's admission in court when it is obtained in violation of the defendant's constitutional rights.[147] For example, if the police enter someone's house without consent or a warrant, as required by the 4th Amendment, and discover incriminating evidence, that evidence cannot be admitted at trial. This is to disincentivize police investigators from violating rights in order to obtain probative information.[148] The Exclusionary Rule and resulting doctrines require that factfinders disregard what may be known to be accurate, relevant, incriminating, and even dispositive evidence. In general, however, there is no *right* to have unlawfully obtained evidence excluded from trial.[149] Notwithstanding certain

---

[147] *Weeks v. U.S.* 232 U.S. 383 (1914), incorporated to the states in *Mapp v. Ohio* 367 U.S. 643 (1961).

[148] Although throughout the 20th Century the Court has given different rationales for the Exclusionary Rule, over time it has all but given up on non-deterrent rationales. In *Nardone v. US*, the Court ruled not only that unlawful phone taps were inadmissible, but also evidence derived from the unlawful tap. (308 US 338, 341)(1939). To do otherwise, the Court held, would be " 'inconsistent with the ethical standards and destructive of personal liberty."(*id*). That was in 1939. In 1961, the court in *Mapp* called the exclusionary rule an "essential part of both the Fourth and Fourteenth Amendments."(at 657). In *Hudson v. Michigan*, the majority writes "the exclusionary rule has never been applied except where its deterrence benefits outweigh its substantial social costs." (*Hudson v. Michigan* 547 US 586, 594 (2006) (internal citations omitted). Note the departure from the lofty language in earlier cases. Indeed, for the majority in *Hudson*, Justice Scalia writes that the Court has revised its view on the expansiveness of Mapp: "we have long since rejected that approach." (at 591). Justice Scalia writes that the "massive remedy" of exclusion "has never been applied except where its deterrence benefits outweigh its substantial social costs."(*Hudson* at 595)(internal citations omitted). In *Hudson,* we see the completion of the Court's transformation from relying on reasons of Fourth Amendment protections or judicial integrity to purely the deterrent benefit of the exclusionary rule.

[149] The rights violation is thought to occur when and only when the unlawful search is conducted, not when that evidence is admitted at trial. In its decision in *Herring v. U.S.*, the Court writes, "…the exclusionary rule is not an individual right…" They continue: "We have repeatedly rejected the argument that exclusion is a necessary consequence of a Fourth Amendment violation." (555 U.S. 135, 141)(2009). But some scholars feel differently. *See,* e.g., Re, "The Due Process Exclusionary Rule." (Locating the right to exclusion in the 5th Amendment's Due Process clause.)

disagreements, we ought to recognize the substantial costs in throwing away or discounting

evidence: more guilty defendants walk free and more plaintiffs are not made whole. Justice

Louis Brandeis famously commented on the Exclusionary Rule: "The criminal is to go free

because the constable has blundered."[150]

Without stating a position on the purpose of the criminal justice system, or of

damage awards, I think I can safely state an uncontroversial but rebuttable presumption:

all relevant evidence should receive its due. Whatever the function of trials, it is

presumptively the case that that function is better-achieved when all relevant evidence is

considered.[151] I don't need an argument for why there are strong reasons in favor of

convicting the guilty, or of declaring a judgment for a plaintiff when they've been harmed.

Disagreement about the function of the trial and the optimal distribution of error will affect

when the presumption is rebutted, but it's important to note the cost of such a rebuttal, as

the Court repeatedly does.[152]

For the kinds of cases described above, most scholars have explained the data and

defended the intuitions. First, they give some account of the difference between statistical

evidence and individualized evidence. Then they explain what (decisive) normative

---

[150] *People v. Defore* 150 N.E. 585, 587 (N.Y. 1926).

[151] Laudan (ms) "The Social Contract and the Rules of Trial: Re-Thinking Procedural Rules." 26-7. Unpublished draft available at http://ssrn.com/abstract=1075403.

[152] *Nardone v. US*, 308 US 388 (1939) at 340: "Any claim for the exclusion of evidence logically relevant in criminal prosecutions is heavily handicapped. It must be justified by an over-riding public policy expressed in the Constitution or the law of the land."; *Nix v. Williams*, 467 U.S. 431 (1984) at 442: "this admittedly drastic and socially costly course is needed to deter police from violations of constitutional and statutory protections."; *Hudson v Michigan*, 547 US 586 (2006) at 591: "The exclusionary rule generates 'substantial social costs[.]"

implications this difference has. Without questioning the distinctions that others have drawn between statistical and individualized evidence, I argue that the normative implications they draw are either misguided, or else, not decisive. That is, the difference between statistical evidence and individualized evidence could be relevant but not decisive in determining whether a factfinder can rest a judgment entirely on statistical evidence. It is good to keep in mind that shortcomings in statistical evidence, whatever they are, are not necessarily decisive against its use.

4.2.2 NON-ACCURACY-BASED EPISTEMIC REASONS FOR THE INSUFFICIENCY OF STATISTICAL EVIDENCE

Philosophers have offered a wealth of views to defend the insufficiency of statistical evidence for conviction or a finding. Mostly, the rationalizations have relied on what I've called non-accuracy-based epistemic considerations. Broadly, philosophers have argued either that statistical evidence lacks a property necessary for knowledge or that it lacks a property necessary for belief.

Conceptual analysis of propositional knowledge is often dated to Socrates' discussion in the *Theaetetus*, in which Socrates gives the justified, true belief account that remained (mostly) popular until Edmond Gettier's influential 1963 paper "Is Justified True Belief Knowledge?" answered the titular question in the negative.[153] So-called Gettier cases are instances in which the subjects have justified true belief and yet the intuition is that they lack knowledge. Prefiguring Gettier, Bertrand Russell offers his famous stopped

---

[153] But see Dutant, "The Legend of the Justified True Belief Analysis." (Giving a revisionist history of the JTB account).

clock case, in which a person walks by a broken clock, which had until recently been

working, at the very time at which the clock is stopped.[154] We're meant to think that, though

the person who walks by is justified in believing the time based on what the clock reads,

and truly believes that it is the time the clock displays, he nevertheless fails to *know* the

time. Gettier, in his paper, offers similar cases, and in the half-century that's followed,

scholars have both offered new cases and tried to determine the missing necessary

condition for knowledge, or else to make more robust the justificatory condition. So-called

Lottery cases, in which one assigns very high probability to their having lost the lottery,

but fails to believe or know it, have occasioned a similar kind of response.[155]

The important question we should ask is whether the kinds of answers to Gettier

and lottery type cases are, if correct in that context, relevant to the legal case. I believe that

they are largely not. Notwithstanding the correct answer to that question, scholars have not

done enough to make the connection explicit. Thus, even if it's the case that Gettier-style

counterexamples show that, safety, say, or adroitness, is necessary for knowledge or belief,

it is a further question what relevance this has for legal fact-finding.

4.2.2.1 Knowledge

Judith Jarvis Thomson argues that in order for evidence to be sufficient for

conviction, there has to be a causal connection between the evidence presented and the

crime. Like knowledge, a verdict should not rest on luck. As Thomson puts it, someone's

---

[154] Russell, *Human Knowledge: Its Scope and Limits*, 170–71.

[155] *See,* e.g., Dogramaci, "A Problem for Rationalist Responses to Skepticism," 359. (Arguing, in part, that the statistical inference in lottery cases does not suffice for knowledge).

"…reason for believing that p is true must ensure, or *guarantee*, that p is true."[156] On one causal account of knowledge, the fact that p has to cause the evidence. For example, in Blue Bus, the bus's redness causally explains the testimony that the bus is red.[157] As Thomson writes, "…this is because the accident-causing cab's actually being red…would causally explain its looking red to that witness."[158] This is what Thomson calls "backward-looking" individualized evidence. In contrast, she writes, if it turns out that Red Cab had given a raucous, booze-soaked party, then this would constitute "forward-looking" individualized evidence that someone from Red Cab had caused the accident.

In contrast to statistical evidence, she writes, individualized evidence "is evidence which is in an appropriate way causally connected with the (putative) fact that the defendant is guilty, and hence (putatively) guarantees the defendant's guilt…"[159] What would suffice for a causal connection, does Thompson think? As Gardiner notes, Thomson's account of the appropriate causal account of knowledge is underspecified, and this is without addressing both the notoriously difficult connection between causation and knowledge and the necessity of knowledge for legal proof.[160]

---

[156] Thomson, "Liability and Individualized Evidence," 208.

[157] See Gardiner, "Legal Burdens of Proof and Statistical Evidence," 4–7. (Offering a treatment of Thomson's argument. Importantly, Gardiner worries Thomson's account of causality is underexplained, and argues that statistical evidence *can* play the relevant causal role Thomson desires. Second, Gardiner worries about Thomson's *guarantee* condition, especially as it would pertain to DNA evidence.)

[158] Thomson, "Liability and Individualized Evidence," 203.

[159] Thomson, 214.

[160] See Gardiner, "Legal Burdens of Proof and Statistical Evidence," 4. Especially fn 13.

In a somewhat lengthy response to the Thomson-type argument, Ferdinand Schoeman points out that the law makes no such distinction between evidence that is and is not causally relevant. Focusing on Thomson's distinction between "internal" and "external evidence," he writes "she does not explain why we should require evidence to be internal before we regard it as reliably probative."[161] That her causality requirement is intuitive but not easy to explain is something Thomson herself recognizes.[162] She essentially compares it to an anti-luck condition for knowledge. She argues that it is unjust for the jury to convict when it is just by luck if the jury gets things right.

But what Thomson does not do, as she says it is difficult to do, is give a principled reason for thinking that this causal connection, even if necessary for knowledge, is necessary for a conviction or finding of liability.

In her book, *Probabilistic Knowledge*, Sarah Moss applies her novel account of knowledge to the statistical evidence paradox. Describing with some approval the work others have done to argue that knowledge is required for holding a defendant responsile, Moss notes, however, that even if the knowledge requirement seems apt for juries in criminal cases, where the standard of proof is very high, it seems inapt in civil cases, where the standard is far lower. As she puts it, "Suppose the criminal standard of proof requires a justified full belief that the defendant is guilty. If this is correct, then what attitude does

---

161 Schoeman, "Statistical vs. Direct Evidence," 190.

162 "If we had individualized evidence…then we would feel considerably less reluctant to impose liability on Red Cab. Why is that? That seems to me a very hard question to answer." Thomson, "Liability and Individualized Evidence," 205. She adds, implicating others: "Friends of the idea that individualized evidence is required for conviction have not really made it clear why this should be thought true." (206).

the civil standard [the applicable standard in Gatecrasher] require?"[163] It can't be full belief;

it has to be something like *more likely than not.* Moss continues:

> …proof by a preponderance of the evidence cannot merely require that the factfinder be justified in having greater than .5 credence that the defendant is liable. … *Gatecrasher* demonstrates that there is still something missing from our explanation of why merely statistical evidence is often insufficient to prove guilty beyond a reasonable doubt, as we have not yet explained why merely statistical evidence is often insufficient for proof by any standard at all.[164]

The advantage of her probabilistic account of knowledge, she says, is that it preserves the

requirement that the jury *know* something, which many others have argued for in the

criminal context, while extending the analysis to civil cases—legal proof, she says, requires

knowledge, no matter the standard of proof.[165] With her new, *probabilistic* account of

knowledge, Moss argues that proof beyond a reasonable doubt requires that the factfinder

"know[] a certain probabilistic content, namely that it is beyond a reasonable doubt that

the defendant is guilty."[166] And, whereas the civil standard can't require that the factfinder

know the defendant is liable, because the standard of proof is merely preponderance of the

evidence (more likely than not), as Moss puts it, "…proof of liability by a preponderance

of the evidence requires that the factfinder know *that the defendant is probably liable*."[167]

As with Thomson, though, Moss' discussion lacks an argument, or even an

explanation for the intuition, that legal proof requires knowledge. *Why* should we think that

---

[163] Moss, *Probabilistic Knowledge*, 207.

[164] Moss, 208.

[165] Moss, 211.

[166] Moss, 201.

[167] Moss, 210. (Emphasis in the original).

defendants have a *right* that the jury *know* that they are probably liable? Such a right has

never before been recognized, or even articulated. One reason to be skeptical of a *right* that

the jury *know* the defendant is probably liable, or any even remotely similar statement, is

that it is absent from a sampling of model civil jury instructions on the preponderance of

the evidence standard from several U.S. circuit courts, on which state jury instructions are

often modeled.[168]

4.2.2.2 Belief

Another strategy in this field comes in response to a focus on formalization and

discussion of degrees of belief, or *credences*, where a credence is the subjective probability

one assigns to the truth of a proposition. There is much debate over the relationship between

---

[168] O'Malley, Grenig, and Lee, *Federal Jury Practice and Instructions*, vol. 3, sec. 101:41.:

3rd Circuit:

> [Plaintiff] has the burden of proving [his/her/its] case by what is called the preponderance of the evidence. That means [plaintiff] has to prove to you, in light of the all the evidence, that what [he/she/it] claims is more likely so than not so. To say it differently: if you were to put the evidence favorable to [plaintiff] and the evidence favorable to [defendant] on opposite sides of the scales, [plaintiff] would have to make the scales tip somewhat on [his/her/its] side. If [plaintiff] fails to meet this burden, the verdict must be for [defendant]. If you find after considering all the evidence that a claim or fact is more likely so than not so, then the claim or fact has been proved by a preponderance of the evidence.

5th: Plaintiff [x] has the burden of proving [his/her] case by a preponderance of the evidence. To establish by a preponderance of the evidence means to prove something is more likely so than not so. If you find that Plaintiff [name] has failed to prove any element of [his/her/its] claim by a preponderance of the evidence, then [he/she/it] may not recover on that claim.

8th: You will have to decide whether certain facts have been proved [by the greater weight of the evidence]. A fact has been proved [by the greater weight of the evidence], if you find that it is more likely true than not true. You decide that by considering all of the evidence and deciding what evidence is more believable.

9th: "When a party has the burden of proving any claim…by a preponderance of the evidence, it means you must be persuaded by the evidence that the claim … is more probably true than not true."

credences and belief—in particular, whether the latter is reduceable to the former. One might think that belief just is credence past a certain threshold—whether a stable threshold or one that changes depending on context.

Lara Buchak weighs in on this debate. She first argues that beliefs do not reduce to credences, and that this is explained, at least in part, by the unique role belief plays in assigning blame. She gives cases that are meant to elicit the intuition that two cases, both of which license the same degree of belief, do not both license *belief*, because in only one of the cases is blame appropriate. In the cases Buchak gives, the only important difference is that the evidence in one case is statistical, and individualized in the other. But, she says, "…what is interesting about statistical evidence is that it is often by itself not enough to produce a belief that *p,* even when [the credence] is very high."[169] Belief and credences are sensitive to different kinds of evidence, she writes. And *blame* is sensitive to belief, but not credence.[170] She articulates the subjective version of the BLAME NORM: "Blame someone if and only if you believe (or know) that she has transgressed…"[171] Because juries are called on not only to determine the facts, but also to "take a stand about whether [the defendant] is guilty," and because (by the BLAME NORM) assigning blame requires belief, *and* because statistical evidence cannot give rise to belief, statistical evidence alone cannot license a jury's verdict.[172]

---

[169] Buchak, "Belief, Credence, and Norms," 292.

[170] Buchak, 296–97.

[171] Buchak, 299.

[172] Buchak, 301.

Andrea Roth makes a similar claim. She writes that the reasonable doubt standard "…requires factfinders to reach an 'actual belief' in, rather than an acknowledgment of a high probability of, the defendant's guilt."[173] Her explanation is largely historical, and draws from the common law notion of "moral certainty," which, at least historically, must be based on testimony and perception, and which the "reasonable doubt" language replaced in the 19th Century.[174]

There are a few ways to dispute Buchak's claims. First, we might bight the bullet and contest her description of the (appropriate) intuitions she has in the cases—perhaps we *would* count ourselves as having a full belief in the statistical cases she gives. Second, we might contest her description of blame as requiring belief and not merely high credence. That is, we might think to ourselves *well, even if I wouldn't have had the belief, my high credence* does *license blame*. But perhaps the most profitable response we could make to Buchak is to question that the role of the trial is the assignment of blame, at least in the way she describes. One might say that if the trial's role is to assign blame, it requires something different from what she requires for interpersonal blame in the cases she presents. And, even if she's right in the criminal context, which it's not clear she's done sufficient work to convince the reader that she is, it's far from clear that the purpose of a civil trial is to assign blame in any sense reminiscent of the reactive attitudes we have

---

173 Roth, "Safety in Numbers? Deciding When DNA Alone Is Enough to Convict," 1159.

174 Roth, 1160.

toward each other.[175] One prominent strain of tort theory—the economic approach—does

not countenance blame, in any moral sense, at all.[176]

4.3 Statistical Evidence as Sufficient for Conviction or a Liability Judgment

4.3.1 THE LEGAL LANDSCAPE

Philosophers who argue for the insufficiency of statistical evidence in the legal

context often overstate the extent to which the courts agree with them. [177] (Though, in their

defense, it is often difficult to track the courts' less than systematic approach to statistical

---

[175] Buchak, "Belief, Credence, and Norms," 304.

[176] Judge Calabresi describes tort law in the following way: "[It is] axiomatic that the principle function of accident law is to reduce the sum of the costs of accidents and the costs of avoiding accidents…" (Cited in Posner, "Guido Calabresi's 'The Costs of Accidents': A Reassessment," 15–16., quoting *The Cost of Accidents* at 26-28.)

[177] *See,* e.g., "It is important to note that the statistical evidence is not *inadmissible;* rather, it is *insufficient* on its own." (Discussing the evidence in the Blue Bus case, in particular. Emphasis in the original). Buchak, "Belief, Credence, and Norms," 290–91.

"In [a version of the Gatecrasher case,] courts will find for the defendant." Blome-Tillmann, "Sensitivity, Causality, and Statistical Evidence in Courts of Law," 103., Adding, "The intuitive distinction between individual and bare statistical evidence can be found in a large number of court judgments and is drawn frequently, with more or less rigour, in the legal and philosophical literature." Blome-Tillmann, 104.

"We would never convict someone of a crime based on statistical evidence alone." Jackson, "Belief, Credence, and Evidence," 5.

"Law courts would not adjudicate in favor of the claimants…" (Giving versions of five canonical cases, including Gatecrasher, Blue Bus, and Prison Yard.) (Gardiner, forthcoming, 3).

Smith 2018 and Dant 1988, respectively, make somewhat weaker claims, which seem a bit misleading nonetheless, (though I don't mean to suggest intentionally so): "Indeed, it seems generally true that courts are reluctant to base affirmative verdicts—verdicts of guilt or liability—on evidence that is purely statistical in nature." Smith, "When Does Evidence Suffice for Conviction?," 1195. Adding, "…courts' general reluctance to rely on purely statistical evidence." (1213, fn 18) but noting, in addition, that "…the legal treatment of statistical evidence has not been entirely consistent." (1195, fn 3).

"Courts and commentators often defend the traditional view that statistical evidence is alone insufficient to support a verdict by appealing to the injustice of imposing liability based on statistical data." Gant, "Gambling on the Truth: The Use of Purely Statistical Evidence as a Basis for Civil Liability," 33.

evidence.[178]) They do not engage, to any significant degree, with the cases in which statistical inference is sufficient for a judgment of liability in the civil context. While it may be true that courts would regard as insufficient the statistical evidence available in the canonical cases presented in the philosophical literature, there are many, many kinds of cases in which the courts *do* permit statistical evidence to be sufficient. For example, in the civil context, statistical evidence is sufficient for judgement in employment discrimination cases and in market share liability cases. Slightly more controversially, statistical evidence is used, decisively, in bail and sentencing determinations, although these are not subject to the same standard of proof considerations as determinations of guilt or liability. As distinct from the use of statistical evidence in trials, the use of so-called "risk-assessment" in bail and sentencing is thought by some to be subject to racial bias,[179] which would constitute both an accuracy-based epistemic reason and a policy reason against its use.

Of course, that judges have held that statistical evidence is sometimes sufficient does not justify the practice, but to the extent it is an endorsed for principled reasons, in a wide variety of cases, despite what is broadly denied or downplayed in the philosophical literature, and to the extent that it constitutes good policy, it does. Below, I highlight several

---

[178] *See,* e.g., the difference in the Supreme Court's holdings in *Tyson Foods, Inc. v. Bouaphakeo* 136 S. Ct. 1036 (2016) and *Wal-Mart Stores, Inc. v. Dukes*, 564 U.S. 388 (2011), as discussed in Bone, "Tyson Foods and the Future of Statistical Adjudication." See also Amit Pundik, who notes that "case law seems to lack a systematic approach to statistical evidence." Pundik, "The Epistemology of Statistical Evidence," 117. But see Koehler, "When Do Courts Think Base Rate Statistics Are Relevant?" (Giving a systematic assessment of when courts seem to approve or disapprove of the sufficiency of statistical evidence).

[179] *See,* e.g., Mayson, "Bias In, Bias Out," 6.; Slobogin, "Principles of Risk Assessment: Sentencing and Policing," 589–93. (Arguing that, properly guided by the relevant principles of fit, validity, and fairness, risk-assessment algorithms are to be preferred over individualized professional judgment).

instances in which the courts have held that statistical evidence is sufficient for liability judgments, and argue that, were we to countenance the arguments made by many philosophers, remedies would be largely unavailable to many different kinds of injured parties.

4.3.1.1 The Criminal Context

As forensic science has become more sophisticated, DNA evidence's use in criminal trials has increased,[180] as have debates about its proper use.[181] In a *pure cold hit*, the DNA match between the defendant and the relevant material is the only evidence. Some criminal defendants have attempted to argue that cold hit DNA evidence cannot be sufficient for conviction because it puts them in a class of suspects rather than uniquely identifying them as the culprit. But, as Andrea Roth puts it, appellate courts have "uniformly rejected" these arguments.[182] This is broadly in line with Jonathan Koehler's descriptive analysis of when courts tend to allow the sufficiency of statistical evidence. DNA cases represent instances where judges tend to allow the sufficiency of statistical evidence because they "rebut the suggestion that the outcome arose by chance." [183]

---

[180] Roth, "Safety in Numbers? Deciding When DNA Alone Is Enough to Convict," 1140–41.

[181] See Wasserman, "Forensic DNA Typing," 349; Semikhodskii, *Dealing with DNA Evidence: A Legal Guide*, 136.

[182] Roth, "Safety in Numbers? Deciding When DNA Alone Is Enough to Convict," 1150. See also *Missouri v. Abdelmalik*, 273 S.W.3d 61 (2008), cited in Di Bello, "Trial by Statistics: Is a High Probability of Guilt Enough to Convict?," 30 fn 46.

[183] Koehler, "When Do Courts Think Base Rate Statistics Are Relevant?," 388–89.

In most cases, the courts have relied on the extremely high probability of the DNA evidence's accuracy.[184] As Roth sees it, the court's willingness to countenance this admittedly and exclusively statistical evidence is because the probabilities of correct DNA matches can be incredibly high—so high as to license "actual belief", or moral certainty. Roth writes, "[a]s the probability of guilt becomes closer and closer to certainty … a statistical statement of the likelihood of guilt may actually be transformed in the jurors' minds from probabilistic evidence to individualized evidence justifying an actual belief in the defendant's guilt."[185] Putting it slightly differently, Roth writes, "…when source probabilities are high enough, they are effectively transformed into statements of certainty rather than probability."[186]

If, as Roth argues, the difference between DNA evidence and other statistical evidence is the extremely high likelihood of a correct match in the DNA evidence, this does not match the explanations for the intuitions expressed about Gatecrasher-style cases in most of the philosophical literature.[187] As one pair of scholars put it,

> [i]t will be interesting to see whether such legal theories will be challenged by the sheer statistical power of the probabilities generated by forensic DNA matches, which, some might say, make the DNA database "a system not of evidence but of proof."[188]

---

[184] Roth, "Safety in Numbers? Deciding When DNA Alone Is Enough to Convict," 1150.

[185] Roth, 1169.

[186] Roth, 1158–59.

[187] Marcello Di Bello is one notable exception. He argues, for accuracy-based epistemic reasons, that we ought not rely on statistics in the Prison Yard-type cases but we may be able to, in DNA evidence cases Di Bello, "Trial by Statistics: Is a High Probability of Guilt Enough to Convict?," 29–32.

[188] Cole and Lynch, "The Social and Legal Construction of Suspects," 51.

But the point of the Prison Yard style cases, I took it, is to show that no matter how many prisoners there are, no matter how confident the jury should be in the defendant's guilt, when that high confidence is based on statistical evidence alone this does not suffice for conviction. Martin Smith's argument for the insufficiency of statistical evidence in the legal domain relies on his "normic support" condition, where "a body of evidence E *normically supports* a proposition P just in case the circumstances in which E is true and P is false would be less normal, in the sense of requiring more explanation, than the circumstances in which E and P are both true."[189] Thus, even where the probability of some event is low it may not call out for much explanation, like, for example, winning the lottery: "The fact that there are 100 tickets in the lottery and only one winner does not normically support the proposition that ticket #72 has lost."[190] Thus, we can't say we *know* that our ticket has lost a large, fair lottery, even though we could know, say, by testimony of a mostly-reliable witness, that the defendant committed the crime. In the Blue Bus case, for example, if the generally reliable witness had got things wrong—if she had hallucinated, or the Yellow Bus company had for some reason painted its busses blue and put "Blue Bus Company" signs on them, then we would expect some sort of explanation. As he puts it,

> It can't 'just happen' that the testimony was wrong. But it *could* just so happen that the bus was not a Blue-bus in spite of the fact that 90% of the buses operating in the area on the day in question were Blue-Bus busses. While this might in a sense be surprising, given the proportions involved, it clearly wouldn't demand any kind of further explanation.[191]

---

[189] Smith, "When Does Evidence Suffice for Conviction?," 1208.

[190] Smith, 1208.

[191] Smith, 1208.

But as he notes, "A normic standard of proof would block pure cold hit DNA convictions," which tend to be allowed by the courts.[192] How should we resolve this discrepancy? "The clash with the normic standard could be portrayed as a reason for being critical of such convictions, but could also be seen as a reason for resisting the standard and seeking an alternative solution to the legal puzzle of statistical evidence."[193] (Though he seems not to see this as a sufficient reason to do so.) Moreover, Smith's account would have trouble countenancing other kinds of cases in which the courts routinely rely on statistical evidence. Would it *call out for explanation* if what seemed to be a pattern of discriminatory behavior in an employment scenario, for example, wasn't intentional? On the normic support view, it seems like it wouldn't.

And, while Sarah Moss discusses the sufficiency of DNA evidence for knowledge, she doesn't quite say whether, on her view of probabilistic knowledge, when (if ever) DNA evidence suffices for knowledge. [194]

One promising response is offered by Enoch and Fisher, which effectively highlights the often-confused way in which scholars discuss "statistical evidence." They describe DNA evidence as "statistical evidence," but argue that one draws not a statistical inference from such data, but rather an inference to the best explanaton. Rather than inferring from the defendant's membership in some class that they are likely guilty, an inference from DNA evidence posits that the best explanation for the DNA match is that

---

[192] Smith, 1214.

[193] Smith, 1214.

[194] Moss, *Probabilistic Knowledge*, 218.

the defendant committed the crime.[195] In this sense, while the evidence does appear statistical in nature, the inference involving it is not properly analogous to the paradoxical cases commonly given. This line of argument is one promising way for those who dislike statistical inference for legal liability to preserve the sufficiency of DNA evidence.[196]

4.3.1.2 The Civil Context

In the civil context, statistical evidence is widely admitted into evidence, and is often sufficient for a judgment for the plaintiff. This is not to say that the Gatecrasher and Blue Bus cases don't have persuasive force. But to use these hypotheticals, or real-world approximations of them, as evidence that we ought not or that the courts *do not* consider similar evidence sufficient, is improper. Many of the civil cases where the courts have held statistical evidence to be sufficient for a finding is when so-called individualized evidence is impossible to gather. Here again, this is broadly in line with Koehler's analysis of courts' inclination toward or against the sufficiency of statistical evidence. Of three common reasons, he finds, that courts permit the evidence to be sufficient, one is where there exists no particularized information.[197] In general, courts have allowed for the sufficiency of statistical evidence when, otherwise, it would be difficult, if not impossible, for genuinely injured parties to collect damages.

---

[195] Enoch and Fisher, "Sense and 'Sensitivity': Epistemic and Instrumental Approaches to Statistical Evidence," sec. III.C.

[196] Thanks to Sinan Dogramaci for bringing this to my attention. As far as I could tell, Enoch and Fisher are the only authors, in either the legal or philosophical literature, to make this distinction.

[197] Koehler, "When Do Courts Think Base Rate Statistics Are Relevant?," 387–88. Adding that market share liability cases "provide another context in which general base rates are sometimes regarded as relevant, admissible, and even dispositive." (at 399).

In toxic torts cases, for example, plaintiffs sue whoever is responsible for exposing them to toxins that they allege are the cause of their medical malady. Because it is often impossible to show that, say, a company's toxic dump *caused* their disease, plaintiffs rely on epidemiological evidence to show causation between the company's dump and their illness by showing an increase in disease relative to the normal incidence of the same disease in the population.

As the court wrote in a representative and much discussed case, in toxic tort cases:

> [because] the chance that there would be particularistic evidence is in most cases quite small, the consequence of retaining the requirement [for particularistic evidence] might be to allow defendants who, it is virtually certain, have injured thousands of people and caused billions of dollars in damages, to escape liability.[198]

Relatedly, in market share liability cases, for example—cases involving an injury from taking a generic drug produced by one of a small number of manufacturers—the courts have ruled in favor of plaintiffs even where they cannot show particularized evidence that the pill that made them ill came from any one manufacturer, in particular. Rather, the manufacturers of the drug are held liable in proportion to their share of the relevant market.[199]

---

[198] *In re Agent Orange Product Liability Litigation*, 597 F.Supp. 740, 836 (1984). For discussion, see Gant, "Gambling on the Truth: The Use of Purely Statistical Evidence as a Basis for Civil Liability," 61–69.

[199] See *Sindel v. Abbot* 1980 611-13; and *Hymowitz* 1989 at 511-512. For discussion, see Kaye, "The Limits of the Preponderance of the Evidence Standard: Justifiably Naked Statistical Evidence and Multiple Causation." (Arguing that, in multi-defendant cases, an interpretation of the preponderance of the evidence standard consistent with the reasoning of e.g., *Sindel*, is appropriate); Koehler, "When Do Courts Think Base Rate Statistics Are Relevant?," 399–400. (Pointing out that *Sindel* is indicative of the courts' willingness to rely on statistical evidence when it would be implausible for plaintiffs to offer individualized evidence); Sara Moss discusses this case very briefly. Moss' view, she writes, can account for the differing intuitions about the ability of statistical evidence to suffice for

Courts have ruled in a similar manner in employment discrimination cases, where, for example, it is alleged that an employer hires too few minority employees given the number of qualified minority employees in the relevant population.[200] And in antitrust law.[201] And securities class action cases.[202] And in class action lawsuits: In *Tyson Foods*, for example, plaintiffs were workers at a meat processing plant who sued when they were not paid for time spent putting on and taking off protective work gear. The court allowed statistical sampling (using *average* times it took workers to dress) in determining whether workers were underpaid.[203]

---

knowledge. But, it's not clear what Moss has to say about these kinds of cases. She writes that her "account of statistical evidence has an unusual strength: it can explain why verdicts against defendants [who are persons,] in *Prison Yard* and *Gatecrasher* seem especially intolerable in comparison with other verdicts [like *Sindel*] that *might or might not be* licensed by statistical evidence." Moss, *Probabilistic Knowledge*, 219. (Emphasis added.)

[200] See Bone, "Tyson Foods and the Future of Statistical Adjudication," 612."…disparate impact in a Title VII [of the Civil Rights Act of 1964, prohibiting employment discrimination on various grounds, including race, sex, and religion] is essentially a statistical concept calling for statistical proof."; Koehler, "When Do Courts Think Base Rate Statistics Are Relevant?," 386.; *Teamsters v. United States*, 431 U.S. 324 (1977) and; *Ricci v. DeStefano*, 557 US. 557 (cited in Pardo and Allen, "Juridical Proof and the Best Explanation," 264, fn 137.); *Hazelwood School District v. US* 433 U.S. 229 (Cited in Koehler, "When Do Courts Think Base Rate Statistics Are Relevant?," 386.)

[201] *See,* Bone, "Tyson Foods and the Future of Statistical Adjudication," 612.: "…statistical modeling is used in antitrust suits to determine damages when it is impossible to know directly what the counterfactual market free from the antitrust violation would have looked like." Adding, in the antitrust case, "statistical evidence is the obvious—and often the only—way to prove the issue and generate a reasonably correct substantive result for each individual case." (id).

[202] *See,* e.g., *Halliburton Co. v. Erica P. John Fund, Inc.* 34 S.Ct. 2398 (Cited in Allen and Pardo, "Relative Plausibility and Its Critics," 24.; Burtis, Gelbach, and Kobayashi, "Error Costs, Legal Standards of Proof, and Statistical Significance," 5.

[203] *Tyson Foods, Inc. v. Bouaphakeo* 136 S.Ct. 1036, 1048. For discussion, see Bone, "Tyson Foods and the Future of Statistical Adjudication"; Allen and Pardo, "Relative Plausibility and Its Critics," 24. Although, there is reason to think that this is a case of statistical *sampling*, and ought to be distinguished from the use of statistical evidence, e.g., base rates, Bone, "Tyson Foods and the Future of Statistical Adjudication," pt. 2.: "[The Supreme Court] treats the case as one involving statistical evidence and employee-specific inferences when it actually involves substituting statistical averages for employee-specific fact finding. This makes it more like a case of statistical adjudication than a case of statistical evidence." Bone, 610., but noting, however, "The distinction

4.3.1.3 Fairness Considerations

The sufficiency of statistical evidence for liability judgments gives some reason to worry about fairness. Not only does it appear unfair to punish a person based on evidence that (admittedly) leaves a large chance of error and has nothing to do with that person, in particular, but it also leads to an odd conclusion: it allows for the double-counting of evidence. Take the Gatecrasher case, for example: If the evidence that 501 attendees went into the rodeo without a ticket is sufficient to convict any one of the 1000 guests, then it is also sufficient to convict *all* of the attendees, 499 of whom *did* buy a ticket. And this seems unfair, if not downright ridiculous. How can one resolve this apparent difficulty with the use of statistical evidence? I see (at least) three ways to respond.

First, one might bite the bullet here and argue that while this seems odd, there is no contradiction or rights violation. Such a result would merely be an instantiation of the preponderance of the evidence standard. Of each individual case, civil cases admit of a high error rate. Why think there's any principled difference between one case with a high chance of error and a large set of cases with a known number of errors? The latter is just what we have now with the body of cases the courts decide. Some objections to this kind of thinking are similar to the legitimacy worries that come up in the legal literature—what I've called *policy* reasons against the use of statistical evidence, and which I've discussed, briefly, above.

---

between statistical adjudication and statistical evidence is not always perfectly clear or precise."
Bone, 613.

Second, one could point out that this already happens in other domains in the law.[204]

In *Bradshaw v. Stumpf*, for example, "the State was on record as maintaining that Stumpf and Wesley should both be executed on the ground that each was the triggerman, when it was undisputed that only one of them could have been."[205] On remand to consider due process violations, the Sixth Circuit held that no rights violation occurred. As Andrew Pollis puts it:

> According to the Sixth Circuit majority, "[a]ll that the prosecution did was to argue for two different inferences from the same, unquestionably complete, evidentiary record." The dissenters had a notably different view, accusing the prosecution of "convenient flip-flopping" that "simply reeks of unfairness." The dissent chastised the "unwavering commitment to a win-at-any-cost callousness that is directly at odds with our solemn oath to preserve and defend the Constitution of the United States." As commentators have noted, the flip-flopping "plainly runs counter to the prosecutor's duty to seek justice." And yet, the legal system is on record as tolerating two death sentences for a crime only one person could have committed.[206]

Anne Poulin argues that this use of evidence is common, and that such use *does* constitute a due process violation.[207] But she notes that courts have considered and rejected a number of challenges to the prosecutors' use of this kind of evidence.[208] Her style of argument does raise some interesting questions about when such a violation occurs. According to Poulin, it occurs at the time the prosecutor asserts the second position.[209] It's

---

[204] This is the general strategy in Pundik, "What Is Wrong with Statistical Evidence? The Attempts to Establish an Epistemic Deficiency."

[205] *Bradshaw v. Stumpf* 545 U.S. 175, 189 (2005). For a description of the case in the popular press, see Armstrong, "Two Murder Convictions for One Fatal Shot."

[206] Pollis, "Trying the Trial," 87–88.

[207] Poulin, "Prosecutorial Inconsistency, Estoppel, and Due Process."

[208] Poulin, 1425 fn 4.

[209] Poulin, 1425.

at least preliminarily suspect to think that a due process violation occurs, as it were, outside the scope of one defendant's trial. That is, how could it be that whether a due process violation occurs at A's trial depends on what happens, say, six months later, in B's trial? If the rights violation occurs, it's odd to think it happens after the trial. But put this oddness aside. Some courts *have* ruled that it is a due process violation for a prosecutor to use one bit of evidence to secure mutually exclusive verdicts.[210] Nunn argues that the use of statistical evidence is a due process violation for a similar but distinct reason: the mere fact that some bit of evidence *could be* used to secure the conviction of any number of people constitutes a due process violation if that evidence is used, even against one person.[211]

As a last response to the problem of double-counting evidence, tort law could move wholesale to a sort of modified proportional liability scheme. In the market share liability cases discussed above, companies were held liable to the extent that their product was represented in the market. At bottom though, the judgement against any company was determined by the likelihood that they caused the harm, and, therefore, the total damage award was capped by the extent of the damage. Part of the worry with relying on statistical evidence in the Gatecrasher case, as discussed above, is that, if statistical evidence is sufficient for conviction, the rodeo owner could recover more than he is owed. If, for example, the tickets each cost $1, and 499 guests paid and 501 guests crashed, the owner

---

[210] Nunn, "The Incompatibility of Due Process and Naked Statistical Evidence," 1418 fn 62–63.

[211] "…if the same naked statistical evidence could be used to convict any randomly selected member of a population, and the simultaneous conviction of the entire population would constitute a due process violation (due to the mutually exclusive nature of the crime), then the conviction of *even one* of those individuals constitutes a due process violation." Nunn, 1427.

is owed $501. But if statistical evidence alone is sufficient for a judgment, he could collect $1 from each of the one thousand guests in 1000 individual trials. This is more than he is owed.

In a modified proportional liability scheme, instead, the owner might be allowed to collect only $.51 from each guest. Thus, the damage award is capped. In a similar manner, in an individual case with statistical evidence, the plaintiff could instead recover damages in proportion to the weight of the evidence. This would require a more precisified analysis of the evidence, and perhaps a more nuanced jury deliberation, but it may be a workaround to some of the problems statistical evidence poses. This is not too far from what Poulin hints at, as a solution. She suggests, as an option to deal with the uncertainty: "…[the prosecutor] may embrace the uncertainty, acknowledge that she cannot prove which of the two pulled the trigger, and adjust her charge and sentencing goals downward."[212]

4.3.2 A PSYCHOLOGICAL EXPLANATION

Both philosophers and legal scholars have offered sophisticated, if greatly varied, arguments for the insufficiency of statistical evidence as proof of guilt or liability. In some cases, the courts have agreed, though in far from a systematic way.

Several psychology studies have found that people are reluctant to make judgments about legal responsibility (civil or criminal) when the evidence is based on naked statistics.[213] This tendency is known as the "Wells Effect" and is named for the author of

---

[212] Poulin, "Prosecutorial Inconsistency, Estoppel, and Due Process," 1424–25.

[213] Wells, "Naked Statistical Evidence of Liability: Is Subjective Probability Enough?"; Wright et al., "Factors Affecting the Use of Naked Statistical Evidence of Liability"; Friedman and Turri, "Is

the first study on this topic. A representative statement from a follow-up study captures the phenomenon: "…people judge that knowledge is less likely to result from probabilistic evidence than from perception, and, moreover, [] people deny that knowledge is gained from probabilistic evidence."[214]

In describing the phenomenon, the experimenters have dismissed explanations of reticence involving an inability to understand the standard of proof, causal relevance, that the evidence is not sufficient to raise subjective probabilities to the necessary level, and fairness. In one study, Niedermeier et al identify what they call the "ease-of-simulation" effect, which is that jurors are more willing to acquit defendants when they can more easily imagine the situation in which the defendant is not guilty—often, but not always, when the evidence is circumstantial.[215]

Probabilistic Evidence a Source of Knowledge?"; Niedermeier, Kerr, and Messé, "Jurors' Use of Naked Statistical Evidence"; Sykes and Johnson, "Probabilistic Evidence Versus the Representation of an Event"; Arkes, Shoots-Reinhard, and Mayes, "Disjunction Between Probability and Verdict in Juror Decision Making."

214 Friedman and Turri, "Is Probabilistic Evidence a Source of Knowledge?," 5–6.

215 "When probabilistic evidence of a defendant's guilt contains information that can be used to build a possible (even if unlikely) scenario in which another party is responsible, jurors will be more reluctant to use that evidence to convict." Adding, however, "…when probabilistic evidence of a defendant's guilt contains little or no information that can be used to construct such an exonerating scenario, jurors will be more willing to rely on that evidence to convict." Niedermeier, Kerr, and Messé, "Jurors' Use of Naked Statistical Evidence," 541–42. See also, for an expanded discussion, Heller, "The Cognitive Psychology of Circumstantial Evidence," 290–98.

Andrea Roth agrees that this can explain often-disparate treatment of DNA evidence.[216] Mike Redmayne, too.[217] And this explanation appears similar to Smith's account of normic support. And, it is supported by Koehler's finding that courts downplay the probative weight of statistical evidence when there is individualized evidence, in addition.[218] When the context does not merely involve the use of base rates or other statistical evidence, but also involves issues of morality, things get even muddier. In a now-famous study, Tetlock et al show that when reasoning about sensitive moral issues, people are (even) less willing to rely on base rates.[219] This finding is echoed in another study about statistical evidence at trials, noting the decidedly moral nature of the trial, especially criminal trials.[220] Indeed, there is some reason to be cautious about relying on intuitions

---

[216] "While juries tend to discount DNA match statistics when they can actually envision examples of other potential suspects in the population who might match, they will treat the match as 'compelling proof' of guilt when they can no longer envision such examples." Roth, "Safety in Numbers? Deciding When DNA Alone Is Enough to Convict," 1168.

[217] Redmayne, "Exploring the Proof Paradoxes," 304:

This 'ease of simulation' explanation for the data also accounts for results in experimental research on DNA evidence. Here it has been found that mathematically equivalent ways of expressing the probative force of a DNA match have different effects on subjects: subjects think guilt more likely when told that 'the probability that the suspect would match the blood drops if he were not the source is 0.1 percent' than when told '1 in 1,000 people in Houston who are not the source would also match the blood drops.' This seems to be because the latter formulation makes the possibility of a match with an innocent person easier to imagine.

[218] Koehler, "When Do Courts Think Base Rate Statistics Are Relevant?" 395.

[219] Tetlock et al., "The Psychology of the Unthinkable," 865–66:

Unparsimonious though it may strike those who aspire to create universal theories of social cognition, the current findings suggest that people place a complex host of superficially ad hoc content constraints on how they execute trade-offs, use base rates, and apply causal schemata to narratives. People who function like intuitive scientists or economists in one setting can be quickly transformed into intuitive moralist-theologians when provoked by assaults on sacred values.

[220] "…what the laws of probability indicate is likely to have occurred—is generally viewed as an unacceptable basis for holding the defendant liable for the actual event." (Discussing Charles

when it comes to statistical evidence: base-rate neglect, for example, has been well-documented.

The psychological studies, however, don't tell us what to do with their findings. What we learn is that the behavior of juries is broadly (but not entirely) in line with philosophical theory: for non-accuracy based epistemic reasons, juries, and to some extent, judges, are less likely to attribute responsibility, or knowledge, when the evidence is statistical. But is there good reason for doing so? As far I could determine, there is no mention of a "Wells *Fallacy*" in the literature, as there is, for example, in descriptions of base rate neglect. [221] The psychological literature merely describes a tendency. Perhaps it offers an explanation for why one philosophical theory has largely tended one way on the question of statistical evidence. But it does not *justify* the theory.

4.4 Conclusion

I've argued that explanations for the insufficiency of statistical evidence, even if they provide solid non-accuracy-based epistemic reasons, do not convincingly show why legal fact-finders should care about these reasons. Indeed, in many civil contexts, where otherwise remedies to injured parties would be impossible to obtain, statistical evidence's sufficiency for liability is necessary. While there may be documented reticence to ascribe

---

Nesson's argument that even juries who ascribe great weight to the evidence nevertheless resist that the defendant committed the crime). Sykes and Johnson, "Probabilistic Evidence Versus the Representation of an Event," 211.

[221] Tversky and Kahneman, "Judgement under Uncertainty: Heuristics and Biases."

blame or liability using statistical inference, it's far from clear that our intuitions here ought to guide findings of legal liability.

References

Adler, Jonathan. "Lying, Deceiving, or Falsely Implicating." *The Journal of Philosophy* 94 (1997): 435–52.

Allen, Ronald J. "On the Significance of Batting Averages and Strikeout Totals: A Clarification of the 'Naked Statistical Evidence' Debate, the Meaning of 'Evidence,' and the Requirement of Proof Beyond a Reasonable Doubt." *Tulane Law Review* 65 (1991): 1093–1110.

Allen, Ronald J, and Michael S Pardo. "Relative Plausibility and Its Critics." *The International Journal of Evidence & Proof*, forthcoming 2019, 1–55.

Allen, Ronald J., and Alex Stein. "Evidence, Probability, and the Burden of Proof." *Arizona Law Review* 55 (2013): 557–602.

Anderson, Elizabeth. *Value in Ethics and Economics*. Cambridge, MA: Harvard University Press, 1993.

Arico, Adam J., and Don Fallis. "Lies, Damned Lies, and Statistics: An Empirical Investigation of the Concept of Lying." *Philosophical Psychology* 26, no. 6 (December 2013): 790–816.

Arkes, Hal R., Brittany Shoots-Reinhard, and Ryan S. Mayes. "Disjunction Between Probability and Verdict in Juror Decision Making: Disjunction Between Probability and Verdict." *Journal of Behavioral Decision Making* 25, no. 3 (July 2012): 276–94.

Armstrong, Ken. "Two Murder Convictions for One Fatal Shot." *The New Yorker*, 2017.

Audi, Robert. "Intuition, Inference, and Rational Disagreement in Ethics." *Ethical Theory and Moral Practice* 11, no. 5 (2008): 475–492.

Ayres, Ian, and Jack M. Balkin. "Legal Entitlements as Auctions: Property Rules, Liability Rules, and Beyond." *Yale Law Journal* 106 (1996): 703.

Bagchi, Aditi. "Distributive Injustice and Private Law." *Hastings Law Journal* 60 (2008): 105.

———. "Distributive Justice and Contract." In *Philosophical Foundations of Contract Law*, 193–211. Oxford: Oxford University Press, 2014.

Ballantyne, Nathan. "Epistemic Trespassing." *Mind*, forthcoming.

112

Bandura, Albert. "Moral Disengagement in the Perpetration of Inhumanities." *Personality and Social Psychology Review* 3, no. 3 (1999): 193–209.

———. "Social Cognitive Theory of Moral Thought and Action." In *Handbook of Moral Behavior and Development: Theory*, edited by William Kurtines and Jacob Gewirtz, 45–103. Hillsdale, N.J.: Erlbaum, 1991.

Barnett, Randy E. "Contract Remedies and Inalienable Rights." *Social Philosophy & Policy* 4 (1986): 179–202.

Benton, Matthew. "Lying, Accuracy and Credence." *Analysis* 78 (2018): 195–198.

———. "Lying, Belief, and Knowledge." In *The Oxford Handbook of Lying*, edited by Jörg Meibauer, 120–33. Oxford: Oxford University Press, 2018.

Blome-Tillmann, Michael. "Sensitivity, Causality, and Statistical Evidence in Courts of Law." *Thought: A Journal of Philosophy* 4, no. 2 (June 2015): 102–12.

Bok, Sissela. *Lying: Moral Choice in Public and Private Life*. New York: Random House, 1978.

Bone, Robert. "Tyson Foods and the Future of Statistical Adjudication." *North Carolina Law Review* 95 (2017): 607–70.

Bowles, Samuel. *The Moral Economy: Why Good Incentives Are No Substitute for Good Citizens*. New Haven: Yale University Press, 2016.

Broome, John, and Adam Morton. "The Value of a Person." *Aristotelian Society Supplementary Volume* 68, no. 1 (1994): 167–98.

Bruckner, Donald. "Strict Vegetarianism Is Immoral." In *The Moral Complexities of Eating Meat*, edited by B. Bramble and B. Fischer, 30–47. New York: Oxford University Press, 2015.

Buchak, Lara. "Belief, Credence, and Norms." *Philosophical Studies* 169, no. 2 (June 2014): 285–311.

Burtis, Michelle, Johnah Gelbach, and Bruce Kobayashi. "Error Costs, Legal Standards of Proof, and Statistical Significance." *Supreme Court Economic Review* 25 (2017): 1.

Capes, Justin A. "Blameworthiness Without Wrongdoing." *Pacific Philosophical Quarterly* 93, no. 3 (2012): 417–437.

Carson, Thomas L. "The Definition of Lying*." *Noûs* 40, no. 2 (June 2006): 284–306.

"Change My View (CMV) • r/Changemyview." Reddit. Accessed January 20, 2018. https://www.reddit.com/r/changemyview/.

Chisholm, Roderick M., and Thomas D. Feehan. "The Intent to Deceive." *The Journal of Philosophy* 74 (1977): 143–59.

Choi, Stephen, Mitu Gulati, and Eric Posner. "Altruism Exchanges and the Kidney Shortage." *Law and Contemporary Problems* 77, no. 3 (2014): 290–322.

Christensen, David. *Putting Logic in Its Place*. Oxford: Oxford University Press, 2004.

Cohen, L.J. *The Probable and the Provable*. Oxford University Press, 1977.

Cole, Simon A., and Michael Lynch. "The Social and Legal Construction of Suspects." *Annual Review of Law and Social Science* 2, no. 1 (December 2006): 39–60.

Coleman, Linda, and Paul Kay. "Prototype Semantics: The English Word Lie." *Language* 57, no. 1 (March 1981): 26.

Colyvan, Mark, Helen M. Regan, and Scott Ferson. "Is It a Crime to Belong to a Reference Class?" *Journal of Political Philosophy* 9, no. 2 (June 2001): 168–81.

Cooter, Robert. "Market Affirmative Action." *San Diego Law Review* 31 (1994): 133.

Cornell, Nicolas. "A Complainant-Oriented Approach to Unconscionability and Contract Law." *University of Pennsylvania Law Review* 164 (2015): 1131.

Dagan, Tsilly, and Talia Fisher. "Rights for Sale." *Minnesota Law Review* 96 (2011): 90.

Davidson, Donald. "What Metaphors Mean." *Critical Inquiry* 5, no. 1, (1978): 31–47.

Demos, Raphael. "Lying to Oneself." *The Journal of Philosophy* 57, no. 18 (September 1, 1960): 588.

Dennett, Daniel. *The Intentional Stance*. Cambridge, MA: MIT Press, 1987.

Di Bello, Marcello. "Trial by Statistics: Is a High Probability of Guilt Enough to Convict?" *Mind*, 2018, 1–40.

Ditmarsch, Hans van, Jan van Eijck, Floor Sietsma, and Yanjing Wang. "On the Logic of Lying." In *Games, Actions and Social Software*, edited by Jan van Eijck and

Rineke Verbrugge, 7010:41–72. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012.

Dogramaci, Sinan. "A Problem for Rationalist Responses to Skepticism." *Philosophical Studies* 168, no. 2 (2014): 355–69.

Dorsey, Dale. "Amorality." *Ethical Theory and Moral Practice* 19, no. 2 (April 1, 2016): 329–42.

Driver, Julia. "The Suberogatory." *Australasian Journal of Philosophy; Sydney, N.S.W.* 70, no. 3 (September 1, 1992): 286–295.

Dutant, Julien. "The Legend of the Justified True Belief Analysis." *Philosophical Perspectives* 29, no. 1 (December 2015): 95–145.

Dworkin, Gerald. "Are These 10 Lies Justified?" *The New York Times*, December 14, 2005. http://opinionator.blogs.nytimes.com/2015/12/14/can-you-justify-these-lies.

Elga, Adam. "Reflection and Disagreement." *Noûs* 41, no. 3 (2007): 478–502.

Elias, Julio Jorge. "The Role of Repugnance in the Development of Markets: The Case of the Market for Transplantable Kidneys." *Social Economics: Current and Emerging Avenues*, 2017, 233.

Enoch, David. "How Is Moral Disagreement a Problem for Realism?" *The Journal of Ethics* 13, no. 1 (2009): 15–50.

Enoch, David, and Talia Fisher. "Sense and 'Sensitivity': Epistemic and Instrumental Approaches to Statistical Evidence." *Stanford Law Review* 67 (2015): 557–610.

Enoch, David, Levi Spectre, and Talia Fisher. "Statistical Evidence, Sensitivity, and the Legal Value of Knowledge." *Philosophy & Public Affairs* 40, no. 3 (June 2012): 197–224.

Fallis, Don. "Are Bald-Faced Lies Deceptive after All?" *Ratio* 28 (2015): 81–96.
———. "Bayesians Don't Tell Knowledge-Lies (and Probably Nobody Else Does Either)," ms. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1666545.
———. "Davidson Was Almost Right about Lying." *Australasian Journal of Philosophy* 91, no. 2 (June 2013): 337–53.

———. "Lying as a Violation of Grice's First Maxim of Quality: Lying as a Violation of Grice's First Maxim of Quality." *Dialectica* 66, no. 4 (December 2012): 563–81.

———. "What Is Lying?" *The Journal of Philosophy* 106 (2009): 29–56.

———. "What Liars Can Tell Us about the Knowledge Norm of Practical Reasoning." *The Southern Journal of Philosophy* 49, no. 4 (December 2011): 347–67.

Faulkner, Paul. "What Is Wrong with Lying?" *Philosophy and Phenomenological Research* 75, no. 3 (2007): 535–57.

Feldman, Richard. "Reasonable Religious Disagreements." In *Philosophers Without Gods: Meditations on Atheism and the Secular*, edited by Louise Antony, 194–214. Oxford University Press, 2007.

Fennell, Lee Anne. "Adjusting Alienability." *Harvard Law Review* 122 (2008): 1403.

Finkelstein, Claire. "Is Risk a Harm?" *University of Pennsylvania Law Review* 151, no. 3 (2003): 963–1001.

Foerster, Thomas. "Moral Offsetting." *The Philosophical Quarterly*, forthcoming.

Frege, Gottlob. "On Sense and Reference." In *Translations from the Philosophical Writings of Gottlob Frege*, edited by Peter Geach and Max Black, 2nd ed. Oxford: Blackwell, 1982.

Fried, Charles. *Right and Wrong*. Cambridge, MA: Harvard University Press, 1978.

———. "The Ambitions of Contract as Promise." In *Philosophical Foundations of Contract Law*, 17–41. Oxford: Oxford University Press, 2014.

Friedman, Ori, and John Turri. "Is Probabilistic Evidence a Source of Knowledge?" *Cognitive Science* 39, no. 5 (July 2015): 1062–80.

Gant, Mary. "Gambling on the Truth: The Use of Purely Statistical Evidence as a Basis for Civil Liability." *Columbia Journal of Law & Social Problems* 22 (1988): 31–70.

Gardiner, Georgi. "In Defence of Reasonable Doubt." *Journal of Applied Philosophy* 34, no. 2 (February 2017): 221–41.

———. "Legal Burdens of Proof and Statistical Evidence." In *Routledge Handbook of Applied Epistemology*, edited by David Coady and James Chase. Routledge, forthcoming.

Green, Stuart P. "Lying, Misleading, and Falsely Denying: How Moral Concepts Inform the Law of Perjury, Fraud, and False Statements." *Hastings Law Journal* 53 (2001): 157–212.

Haidt, Jonathan. *The Righteous Mind: Why Good People Are Divided by Politics and Religion*. New York: Vintage Books, 2012.

Harman, Elizabeth. "Morally Permissible Moral Mistakes." *Ethics* 126, no. 2 (2016): 366–393.

Hasen, Richard L. "Vote Buying." *California Law Review* 88 (2000): 1323.

Heller, Kevin. "The Cognitive Psychology of Circumstantial Evidence." *Michigan Law Review* 105 (2006): 241–305.

Horowitz, Sophie. "Immoderately Rational." *Philosophical Studies* 167, no. 1 (January 2014): 41–56.

Horton, Joe. "The All or Nothing Problem." *The Journal of Philosophy* 114 (2017): 94–104.

Isenberg, Arnold. "Deontology and the Ethics of Lying." *Philosophy and Phenomenological Research* 24, no. 4 (June 1964): 463.

Jackson, Elizabeth. "Belief, Credence, and Evidence." *Synthese*, forthcoming, 1–20.

Jeffrey, Richard. "Dracula Meets Wolfman: Acceptance vs. Partial Belief." In *In Induction, Acceptance, and Rational Belief*, edited by M. Swain, 157–85. Dordrecht: Reidel Publishing Company, 1970.

Joyce, James M. "A Nonpragmatic Vindication of Probabilism." *Philosophy of Science* 65, no. 4 (December 1998): 575–603.

———. "Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief." In *Degrees of Belief*, edited by Franz Huber and Christoph Schmidt-Petri, 263–97. Dordrecht: Springer Netherlands, 2009.

Kant, Immanuel. *The Groundwork of the Metaphysics of Morals*, 1785.

Katz, Larissa. "Spite and Extortion: A Jurisdictional Principle of Abuse of Property Right." *Yale Law Journal* 122 (2012): 1444.

Kawall, Jason. "Friendship and Epistemic Norms." *Philosophical Studies* 165, no. 2 (2013): 349–70.

Kaye, David. "The Limits of the Preponderance of the Evidence Standard: Justifiably Naked Statistical Evidence and Multiple Causation." *American Bar Foundation Research Journal* 7, no. 2 (1982): 487–516.

———. "The Paradox of the Gatecrasher and Other Stories." *Arizona State Law Journal*, 1979, 101–9.

Keiser, Jessica. "Bald-Faced Lies: How to Make a Move in a Language Game without Making a Move in a Conversation." *Philosophical Studies* 173, no. 2 (February 2016): 461–77.

Keller, Simon. "Friendship and Belief." *Philosophical Papers* 33, no. 3 (November 2004): 329–51.

Kelly, Daniel B. "Strategic Spillovers." *Columbia Law Review* 111 (2011): 1641.

Kelly, Thomas. "Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research* 66 (2003): 612–40.

Koehler, Jonathan J. "When Do Courts Think Base Rate Statistics Are Relevant?" *Jurimetrics* 42, no. 4 (2002): 373–402.

Koehler, Jonathan J, and Daniel Shaviro. "Veridical Verdicts: Increasing Verdict Accuracy Through the Use of Overtly Probabilistic Evidence and Methods." *Cornell Law Review* 75 (1990): 247–79.

Konek, Jason, and Benjamin A Levinstein. "The Foundations of Epistemic Decision Theory." *Mind* 128, no. 509 (January 1, 2019): 69–107.

Kornhauser, Lewis A., and Lawrence G. Sager. "Unpacking the Court." *Yale Law Journal* 96 (1986): 82.

Krauss, Sam Fox. "Lying, risk and accuracy." *Analysis* 77(4) (2017): 726–34.

Kuflik, Arthur. "The Inalienability of Autonomy." *Philosophy & Public Affairs*, 1984, 271–298.

Kupfer, Joseph. "The Moral Presumption against Lying." *The Review of Metaphysics* 36, no. 1 (1982): 103–26.

Lackey, Jennifer. "Lies and Deception: An Unhappy Divorce." *Analysis* 73, no. 2 (2013): 236–48.

Laudan, Larry. *The Law's Flaws*. College Publications, 2016.

———. *Truth, Error, and Criminal Law: An Essay in Legal Epistemology*. Cambridge, MA: Cambridge University Press, 2006.

Leland, Patrick R. "Rational Responsibility and the Assertoric Character of Bald-Faced Lies." *Analysis* 75 (2015): 550–54.

Levmore, Saul. "Voting with Intensity." *Stanford Law Review* 53 (2000): 111.

Liberto, Hallie Rose. "Denying the Suberogatory." *Philosophia* 40, no. 2 (2012): 395–402.

List, Christian. "The Theory of Judgment Aggregation: An Introductory Review." *Synthese* 187, no. 1 (2012): 179–207.

List, Christian, and Ben Polak. "Introduction to Judgment Aggregation." *Journal of Economic Theory* 145, no. 2 (2010): 441–466.

Mahon, James Edward, and Edward N. Zalta. "The Definition of Lying and Deception." In *The Stanford Encyclopedia of Philosophy*, 2015. http://plato.stanford.edu/archives/spr2016/entries/lying-definition/.

Mannison, Don S. "Lying and Lies." *Australasian Journal of Philosophy* 47 (1969): 132–44.

Marsili, Neri. "Lying as a Scalar Phenomenon: Insincerity Along the Certainty-Uncertainty Continuum." In *In Certainty-Uncertainty–and the Attitudinal Space in Between*, edited by S. Cantarini, W. Abraham, and E. Leiss, 153–73. Philadelphia: John Benjamins, 2014.

Mayer, Jane. *Dark Money*. New York: Doubleday, 2016.

Mayson, Sandra. "Bias In, Bias Out." *Yale Law Journal* 128 (forthcoming 2019).

McCullough, Colleen. "Unconscionability as a Coherent Legal Concept." *University of Pennsylvania Law Review* 164 (2015): 779.

Meibauer, Jörg. "Lying and Falsely Implicating." *Journal of Pragmatics* 37 (2005): 1373–99.

Mercy For Animals. "Farm to Fridge Inspires Students to Boycott Factory Farm Cruelty," http://www.mercyforanimals.org/farm-to-fridge-inspires-students-to-boycott-factory-farm-cruelty;

Meyers, C.D. "Why It Is Morally Good to Eat (Certain Kinds of) Meat." *Southwest Philosophy Review* 29, no. 1 (2013): 119–126.

Michaelson, Eliot. "The Lying Test." *Mind & Language* 31, no. 4 (September 2016): 470–99.

Montaigne, Michel de. *The Essays*. Translated by John Florio. New York: G.P. Putnam's Sons, 1595.

Moss, Sarah. *Probabilistic Knowledge*. Oxford: Oxford University Press, 2018.

Nadelhoffer, Thomas. "The Butler Problem Revisited." *Analysis* 64, no. 3 (2004): 277–84.

National Organ Transplant Act, Pub. L. No. 98–507 (1984).

Nesson, Charles. "The Evidence or the Event? On Judicial Proof and the Acceptability of Verdicts." *Harvard Law Review* 98, no. 7 (May 1985): 1357.

Nesson, Charles R. "Reasonable Doubt and Permissive Inferences: The Value of Complexity." *Harvard Law Review* 92, no. 6 (April 1979): 1187.

Niedermeier, Keith E., Norbert L. Kerr, and Lawrence A. Messé. "Jurors' Use of Naked Statistical Evidence: Exploring Bases and Implications of the Wells Effect." *Journal of Personality and Social Psychology* 76, no. 4 (1999): 533–42.

Nunn, G. Alexander. "The Incompatibility of Due Process and Naked Statistical Evidence." *Vanderbilt Law Review* 68 (2015): 1407–33.

O'Malley, Kevin, Jay Grenig, and William Lee. *Federal Jury Practice and Instructions*. 6th ed. Vol. 3. Thompson West, 2019.

Oman, Nathan B. *The Dignity of Commerce*. Chicago: University of Chicago Press, 2016.

Ord, Toby. "Moral Trade." *Ethics* 118 (2015): 118–38.

Pace, Michael. "The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' 'Will To Believe.'" *Noûs* 45, no. 2 (2011): 239–68.

Pardo, Michael S. "The Paradoxes of Legal Proof: A Critical Guide." *Boston University Law Review* 99 (2019): 223–90.

Pardo, Michael S., and Ronald J. Allen. "Juridical Proof and the Best Explanation." *Law and Philosophy* 27, no. 3 (May 2008): 223–68.

Perillo, Joseph M. "Abuse of Rights: A Pervasive Legal Concept." *Pacific Law Journal* 27 (1995): 37.

Placani, Adriana. "When the Risk of Harm Harms." *Law and Philosophy* 36, no. 1 (2017): 77–100.

Pollis, Andrew. "Trying the Trial." *George Washington Law Review* 84 (2016): 55–120.

Porat, Ariel, and Alex Stein. *Tort Liability under Uncertainty*. Oxford: Oxford University Press, 2001.

Posner, Eric A., and Cass R. Sunstein. "Moral Commitments in Cost-Benefit Analysis," 2017.

Posner, Eric A., and E. Glen Weyl. "Quadratic Voting and the Public Good: Introduction." *Public Choice* 172, no. 1–2 (2017): 1–22.

———. "Voting Squared: Quadratic Voting in Democratic Politics." *Vanderbilt Law Review* 68 (2015): 441.

Posner, Richard A. "Guido Calabresi's 'The Costs of Accidents': A Reassessment." *Maryland Law Review* 64 (2005): 12–23.

Poulin, Anne Bowen. "Prosecutorial Inconsistency, Estoppel, and Due Process: Making the Prosecution Get Its Story Straight." *California Law Review* 89, no. 5 (October 2001): 1423.

Primoratz, Igor. "Lying and the 'Method of Ethics.'" *International Studies in Philosophy* 16 (1984): 35–57.

Pritchard, Duncan. "Risk." *Metaphilosophy* 46, no. 3 (July 2015): 436–61.

Pundik, Amit. "The Epistemology of Statistical Evidence." *The International Journal of Evidence & Proof* 15, no. 2 (April 2011): 117–43.
———. "What Is Wrong with Statistical Evidence? The Attempts to Establish an Epistemic Deficiency." *Civil Justice Quarterly* 27 (2008): 461–93.

Quong, Jonathan. "Cultural Exemptions, Expensive Tastes, and Equal Opportunities." *Journal of Applied Philosophy* 23 (2006): 55–73.

Radin, Margaret Jane. *Contested Commodities*. Cambridge, MA: Harvard University Press, 2000.

———. "Market-Inalienability." *Harvard Law Review* 100, (1987): 1849–1937.

Ramsey, F.P. *Truth and Probability*. Cambridge: Cambridge University Press, 1926.

Re, Richard. "The Due Process Exclusionary Rule." *Harvard Law Review* 127, no. 7 (2014): 1885–1966.

Redmayne, Mike. "Exploring the Proof Paradoxes." *Legal Theory* 14, no. 04 (December 2008): 281.

Rees, C. F. "Better Lie!" *Analysis* 74, no. 1 (January 1, 2014): 59–64.

Regan, Tom. *The Case for Animal Rights*. Berkeley: University of California Press, 1983.

Robinson, Glen O. "Risk, Causation, and Harm." In *Liability and Responsibility: Essays in Law and Morals*, edited by R. Frey and C. Morris, 317–50. Cambridge: Cambridge University Press, 1991.

Rose-Ackerman, Susan. "Inalienability and the Theory of Property Rights." *Columbia Law Review* 85 (1985): 931–69.

Roth, Alvin E. "Repugnance as a Constraint on Markets." *The Journal of Economic Perspectives* 21, no. 3 (2007): 37–58.

———. "What Have We Learned from Market Design?" *Innovation Policy and the Economy* 9, no. 1 (January 2009): 79–112.

Roth, Andrea. "Safety in Numbers? Deciding When DNA Alone Is Enough to Convict." *New York University Law Review* 85 (2010): 1130.

Russell, Bertrand. *Human Knowledge: Its Scope and Limits*. Oxon: Routledge, 1948.

Sachs, Stephen E. "Saving Toby: Extortion, Blackmail, and the Right To Destroy." *Yale Law & Policy Review* 24 (2006): 251.

Sakama, Chiaki, Martin Caminada, and Andreas Herzig. "A Logical Account of Lying." In *Logics in Artificial Intelligence*, edited by T Janhunen and I. Niemela. Berlin: Springer-Verlag, 2010.

Sandel, Michael. "How Markets Crowd Out Morals." *Boston Review*, 2012.

———. *What Money Can't Buy: The Moral Limits of Markets*. New York: Farrar, Straus and Giroux, 2009.

Satz, Debra. *Why Some Things Should Not Be for Sale: The Moral Limits of Markets*. Oxford Political Philosophy. New York: Oxford University Press, 2010.

Saul, Jennifer. "Just Go Ahead and Lie." *Analysis* 72, no. 1 (2012): 3–9.

Scanlon, Thomas. "A Theory of Freedom of Expression." *Philosophy & Public Affairs* 1 (1972): 204–226.

Schauer, Frederick. *Profiles, Probabilities, and Stereotypes*. Harvard University Press, 2003.

Schoeman, Ferdinand. "Statistical vs. Direct Evidence." *Noûs* 21, no. 2 (1987): 179.

Schoenfield, Miriam. "An Accuracy Based Approach to Higher Order Evidence." *Philosophy and Phenomenological Research* 96, no. 3 (2018): 690–715.

———. "Bridging Rationality and Accuracy." *The Journal of Philosophy* 112 (2015): 633–57.
———. "Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief." *Noûs* 47, no. 1 (2013): 193–218.

Semikhodskii, Andrei. *Dealing with DNA Evidence: A Legal Guide*. Abingdon, UK: Routledge-Cavendish, 2007.

Shaviro, Daniel. "Statistical-Probability Evidence and the Appearance of Justice." *Harvard Law Review* 103, no. 2 (December 1989): 530.

Shibles, Warren. *Lying: A Critical Analysis*. Whitewater, WI: The Language Press, 1985.

Shiffrin, Seana Valentine. "Paternalism, Unconscionability Doctrine, and Accommodation." *Philosophy & Public Affairs* 29, no. 3 (2000): 205–50.

———. "The Divergence of Contract and Promise." *Harvard Law Review* 120 (2007): 708.

Siegler, Frederick A. "Lying." *American Philosophical Quarterly* 3, no. 2 (1966): 128–36.

Singer, Peter. *Animal Liberation*. 4th ed. New York: HarperCollins Publishers, 2009.

———. "Famine, Affluence, and Morality." *Philosophy & Public Affairs* 1 (1972): 229–243.

Slobogin, Christopher. "Principles of Risk Assessment: Sentencing and Policing." *Ohio State Journal of Criminal Law* 15 (2018): 583.

Smith, Martin. "When Does Evidence Suffice for Conviction?" *Mind* 127, no. 508 (October 1, 2018): 1193–1218.

Sorensen, Roy. "Bald-Faced Lies! Lying Without the Intent to Deceive." *Pacific Philosophical Quarterly* 88, no. 2 (June 2007): 251–64.

———. "Knowledge-Lies." *Analysis* 70, no. 4 (October 1, 2010): 608–15.

———. "Lying with Conditionals." *The Philosophical Quarterly* 62, no. 249 (2012): 820–32.

Staffel, Julia. "Reply to Roy Sorensen, 'Knowledge-Lies.'" *Analysis* 71 (2011): 300–302.

Stalnaker, Robert. *Inquiry*. Cambridge, MA: MIT Press, 1984.
Stilz, Anna. "Is the Free Market Fair?" *Critical Review* 26, no. 3–4 (2014): 423–38.

Stokke, Andreas. "Lying and Asserting." *The Journal of Philosophy* 110 (2013): 33–60.

———. "Lying and Misleading in Discourse." *Philosophical Review* 125, no. 1 (2016): 83–134.

Strahilevitz, Lior Jacob. "The Right to Destroy." *Yale Law Journal* 114 (2005): 781.

Stroud, Sarah. "Epistemic Partiality in Friendship." *Ethics* 116, no. 3 (April 2006): 498–524.

Sullivan, Sean P. "A Likelihood Story: The Theory of Legal Fact-Finding." *University of Colorado Law Review* 90, no. 1 (forthcoming 2019).

Sunstein, Cass R. "Incommensurability and Valuation in Law." *Michigan Law Review* 92, no. 4 (1994): 779–861.

Sykes, Deanna L., and Joel T. Johnson. "Probabilistic Evidence Versus the Representation of an Event: The Curious Case of Mrs. Prob's Dog." *Basic and Applied Social Psychology* 21, no. 3 (1999): 199–212.

Tetlock, Philip E., Orie V. Kristel, S. Beth Elson, Melanie C. Green, and Jennifer S. Lerner. "The Psychology of the Unthinkable: Taboo Trade-Offs, Forbidden Base Rates, and Heretical Counterfactuals." *Journal of Personality and Social Psychology* 78, no. 5 (2000): 853–70.

Thaler, Richard, and Cass R. Sunstein. *Nudge: Improving Decisions about Health, Wealth and Happiness*. New York: Yale University Press, 2008.

Thomson, Judith Jarvis. "Liability and Individualized Evidence." *Law and Contemporary Problems* 49, no. 3 (1986): 199.

Tribe, Laurence H. "Trial by Mathematics: Precision and Ritual in the Legal Process." *Harvard Law Review* 84, no. 6 (1971): 1329.

Tversky, Amos, and Daniel Kahneman. "Judgement under Uncertainty: Heuristics and Biases." *Science* 185 (1974): 1124–31.

UCSF. "The Kidney Project," https://pharm.ucsf.edu/kidney/need/statistics.

Urmson, J.O. "Saints and Heroes." In *Essays in Moral Philosophy*, edited by About us, Centre for Digital Philosophy, UWO, Phiosophy Documentation Center, Institute of Philosophy, London, Editorial team, General Editors:, et al., 198–216. Seattle: University of Washington Press, 1958.

U.S. Bureau of Labor Statistics. "Consumer Expenditure Survey," 2016.

Vavova, Katia. "Moral Disagreement and Moral Skepticism." *Philosophical Perspectives* 28, no. 1 (2014): 302–333.

Waismann, Friedrich. *The Principles of Linguistic Philosophy*. Glasgow: The University Press, 1965.

Wasserman, David. "Forensic DNA Typing." In *A Companion to Genethics*, edited by Justine Burley and John Harris, 349–63. Oxford, UK: Blackwell Publishing Ltd, 2004.

Wasserman, David T. "The Morality of Statistical Proof and the Risk of Mistaken Liability." *Cardozo Law Review* 13 (1991): 935–76.

Webber, Jonathan. "Liar!" *Analysis* 73, no. 4 (2013): 651–59.

Wells, Gary L. "Naked Statistical Evidence of Liability: Is Subjective Probability Enough?" *Journal of Personality and Social Psychology* 62, no. 5 (1992): 739–52.

Wiles, Anne. "Lying: Its Inconstant Value." *The Southern Journal of Philosophy* 26 (1988): 275–84.

Williams, Bernard. *Truth and Truthfulness: An Essay in Genealogy*. Princeton: Princeton University Press, 2002.

Williston on Contracts. 4th ed., 2017.

Woodle, E. S., J. A. Daller, M. Aeder, R. Shapiro, T. Sandholm, V. Casingal, D. Goldfarb, et al. "Ethical Considerations for Participation of Nondirected Living Donors in Kidney Exchange Programs: Ethical Considerations for Participation of Nondirected Living Donors in Kidney Exchange Programs." *American Journal of Transplantation* 10, no. 6 (2010): 1460–67.

Wright, Edward, Lora MacEachern, Elaine Stoffer, and Nancy MacDonald. "Factors Affecting the Use of Naked Statistical Evidence of Liability." *The Journal of Social Psychology* 136, no. 6 (1996): 677–88.