

UNIVERSITY OF CALIFORNIA

Los Angeles

A Deliberative Conception of
Moral Complicity

A dissertation submitted in partial satisfaction of the
Requirements for the degree Doctor of Philosophy
in Philosophy

by

Jenna Lyn-Adsit Donohue

2021

© Copyright by

Jenna Lyn-Adsit Donohue

2021

ABSTRACT OF THE DISSERTATION

A Deliberative Conception of Moral Complicity

by

Jenna Lyn-Adsit Donohue

Doctor of Philosophy in Philosophy

University of California, Los Angeles, 2021

Professor Alexander Jacob Julius, Chair

Suppose Agatha tells her roommate Brianna that she plans to cheat on her upcoming midterm exam, going so far as to explain her plan in detail. Agatha explains that many of their fellow students cheat and not to do so is to put oneself at a disadvantage. Brianna nods along as Agatha shares her plan. Though Brianna hopes that Agatha will not cheat, she does not tell Agatha of this hope nor does she raise any of the reasons that cheating is wrong. As it happens, Agatha would have cheated even if Brianna had raised her concerns. It seems to me that Brianna is complicit in Agatha's cheating, and that her failure to responsibly talk through with Agatha what Agatha what should do is what grounds her complicity. In this dissertation, I develop and defend this pretheoretical intuition—that failing to discuss with others what's to be done and why can make you complicit in their wrongful actions. According to what I call the *deliberative view* of complicity, an agent B's failure in her deliberative duty regarding another agent A's φ -ing is sufficient for and explains B's complicity in A's φ -ing, if A φ 's. Deliberative duties are obligations we have with respect to the practical deliberations of other agents. Examples include speaking sincerely, meeting a reasonable person

standard of accuracy in one's speech, and exercising due care that others' false beliefs are not formed nor confirmed on the basis of one's conduct. Violating these duties opens an agent up to moral complicity when others go on to act wrongly. The deliberative view can explain cases that other philosophical accounts of complicity cannot, including silence as complicity, complicity without difference-making, and complicity without intentional participation.

The deliberative view can also explain cases of complicity that do not involve speech. Suppose Betsy drives Adam to a bank so that Adam can rob it. Betsy intuitively seems complicit in Adam's bank-robbing, but this can seem like a difficult case for the deliberative view. Even if Betsy had told Adam that robbing a bank is a wrongful action, it still seems like his *driving* Adam to the bank makes him complicit in Adam's bank-robbing. To help explain this kind of case, I defend a theory of action according to which our conduct expresses our stance on moral issues. When Betsy drives Adam to the bank, he *acts-as-though* Adam's bank-robbing is permissible and so *expresses* that Adam's bank-robbing is permissible. In so doing, he violates his deliberative obligation with respect to Adam's bank-robbing and so opens himself up to complicity in the robbery. More generally, when an action *acts-as-though* an agent A's (wrongful) φ -ing is morally permissible, she opens herself up to complicity in A's φ -ing *because* she has violated her deliberative duty in expressing its permissibility.

A deliberative account of moral complicity can help us to reimagine our understanding of our moral impact on the world and so help us make progress on difficult questions of our complicity in an increasingly global society. By helping others to see the moral truth, I can not only avoid complicity in wrongdoing but also support others who are aiming to act rightly. Through my speech and action, I express my stance on moral issues, and ensuring that those stances are both correct and in line with my actual thinking is important not only to my own moral flourishing but also the moral success of those around me.

The dissertation of Jenna Lyn-Adsit Donohue is approved.

Seana Shiffrin

Barbara Herman

Katrina Elliott

Pamela Hieronymi

Alexander Jacob Julius, Committee Chair

University of California, Los Angeles

2021

For Kevin

without whom this wouldn't have been possible

And for everyone who believed in me

without whom I wouldn't have dreamed I could do this

TABLE OF CONTENTS

<i>Acknowledgements</i>	vii
<i>Vita</i>	ix
0. Introduction	1
1. Chapter 1: The Need for a Deliberative Conception of Moral Complicity	7
2. Chapter 2: Deliberative Duties and Our Default Entitlement to Adopt a Deliberative Stance Toward Others	29
3. Chapter 3: Silence as Complicity and Action as Silence	45
4. Chapter 4: The Scope and Explanation of Deliberative Complicity	62
5. Conclusion	81
6. Bibliography	83

Acknowledgements

I have been thinking about the issues in this dissertation for a long time, from high school to Notre Dame to graduate school. As such, I have benefitted enormously from the kind time and attention of many friends and mentors and am very grateful to each of them.

I am grateful to my committee. A.J. Julius has always pushed me to the next level in my thinking, encouraging me to dig further and to be willing to extend myself a bit more than I would without him. I will cherish long talks in the beautiful LA weather outside LuValle and forever be proud of making him laugh so hard he cried while we discussed the trolley problem in first year seminar. Pamela Hieronymi has offered the kind of support and mentorship that a younger version of myself used to dream about. I am so lucky to have benefited from her careful criticism, her stimulating conversation, and her patient help. I am especially grateful to call her not only an advisor but a friend. Seana Shiffrin has been a source of inspiration and strength. She has constantly pushed me to be a clearer writer, a better thinker, and a more careful reader. She is often the voice in my head, encouraging me to be both more patient with myself and to believe that I am capable of more than I may think. Barbara Herman has always offered a fresh perspective: a new angle, a different example, a distinct way of thinking about a problem. Every conversation with her helped my thinking grow as she pushed me to think a few more steps into the ideas than I would go willingly on my own. I may be most proud of the joy my loud laugh brings her. Katie Elliott has walked with me through good times and hard ones, offering a laugh or a sage piece of wisdom (and often both). I am so grateful to have had her perspective on both my philosophical work and academic career. I hope for many more opportunities to learn from and talk with her.

I am grateful to all of the graduate students at UCLA during my time: we have a special community that pursues truth together. My work and my thinking have benefitted enormously from time spent with all of you. I am particularly grateful to Gabe Dupre, Ashley Feinsinger, Brian

Hutler, Gabbrielle Johnson, Amber Kavka-Warren, Brad McHose, Alexi Patsaouras, Ally Peabody-Smith, Ayana Samuel, Jordan Wallace-Wolf, and all of the members of the Ethics Writing Seminar during my time participating. A special thank-you to Ally, Ayana, and Amber for their mental, emotional, and professional support before, during, and after the job market, to Gabby for serving as a mentor, in both formal and informal capacities throughout my entire graduate career, and to Jordan for many long and patient conversations in our shared office.

I am grateful to many UCLA faculty outside my committee. Our department is a special place where anyone was willing to talk with me about my work. I have benefited in particular from discussions with Josh Armstrong, Daniela Dover, Calvin Normore, Larry Sager, and Rebecca Stone.

Sean Greenberg has offered passionate criticism and lively discussion, and we have shared stimulating and serious disagreement. I will miss our West LA walks, chats, and coffee houses.

At Notre Dame, Timothy Bays, Samuel Newlands, and Peter van Inwagen encouraged me to recognize my passion for philosophy and keep following where it led.

Without my friends and family, I wouldn't be where I am today. Special thanks to Allie Schmitt for always believing in me and Annie Geffroy for always bringing a smile to my face. Annie Gutierrez-DuBois has taken stellar care of my children while I pursued the tricky balance of being both a mother and a scholar, and Sandy, Bill, Michael, and Kira Adsit have always offered their love.

Finally, Kevin—you found new ways to create space for me to work: from making dinner to taking four kids to Costco on your own, and this dissertation wouldn't have gotten done without your support. Thank you. John Paul, James, Lucía, and Peter: I am grateful for your smiles and so proud to be your mother.

Vita

Education

University of California, Los Angeles, M.A., philosophy	2015
University of Notre Dame, M.Ed.	2012
University of Notre Dame, B.A., philosophy (with honors) and B.S., mathematics <i>magna cum laude</i>	2010

Presentations (* indicates refereed)

“Acting as Though: Demonstrating through Conduct that φ is Morally Permissible”:
Central APA Symposium, 2021*
UCLA Prospective Week Presentation, 2019
Ethics Writing Seminar, UCLA, 2018

“Silence as Moral Complicity”: Research Talk, University of North Carolina Wilmington, 2021
Practice Job Talk, UCLA, 2021
Ethics Workshop, UCLA, 2020

“Adopting Active-Learning Strategies for Remote Instruction”:
American Association of Philosophy Teachers 2020 Conference: *Anti-Racist and Pandemic Pedagogies, Pandemic Pedagogies Panel**

“The Need for Deliberative Complicity”: Job Market Seminar, UCLA, 2020

“The Scope of Deliberative Complicity”: UCLA Albritton Society, 2020

“Expressive Action and our Duty of Care to the Mental Contents of Others”:
Ethics Workshop, UCLA, 2019

“What Makes Manipulation Distinctively Wrong?”: Central APA Colloquium, 2019*
UCLA Albritton Society, 2019

“When is it Unreasonable to be Unsurprised?”: Pacific APA Colloquium Speaker, 2018*
Grad Slam Competition, UCLA, 2016
UCLA Albritton Society, 2014

“Is it Possible to Facilitate Autonomy While Avoiding Complicity with Wrongdoing?”:
UCLA Albritton Society, 2017

“Referral as Complicity”: Ethics Workshop, UCLA, 2017
UCLA Proposition Series, 2016

“Kant on the Permissibility of Breastmilk Sale”: UCLA Albritton Society, 2016

Introduction

Suppose Agatha tells her roommate Brianna that she plans to cheat on her upcoming midterm exam, going so far as to explain her plan in detail. Agatha explains that many of their fellow students cheat and not to do so is to put oneself at a disadvantage. Brianna nods along as Agatha shares her plan. Though Brianna hopes that Agatha will not cheat, she does not tell Agatha of this hope nor does she raise any of the reasons that cheating is wrong. As it happens, Agatha would have cheated even if Brianna had raised her concerns. It seems to me that Brianna is morally complicit in Agatha's cheating, and that her failure to responsibly talk through with Agatha what Agatha should do is what grounds her complicity.

In this dissertation, I develop and defend this pretheoretical intuition—that failing to discuss with others what's to be done and why can make you morally complicit in their wrongful actions. According to what I call the *deliberative view* of moral complicity, an agent B's failure in her deliberative duty regarding another agent's wrongful φ -ing is sufficient for and explains B's moral complicity in A's φ -ing, if A φ 's. Deliberative duties are obligations we have with respect to the practical deliberation of other agents. Examples include speaking sincerely, meeting a reasonable person standard of accuracy in one's speech, exercising due care that others' false beliefs are not formed nor confirmed on the basis of one's conduct, and a willingness to engage with other agents in a sincere way as they deliberate about what's to be done. Violating these duties opens an agent up to moral complicity when others go on to act wrongly.

What is at stake when we call someone morally complicit? In other words, what does it amount to when I claim that B's deliberative duty failure with respect to A's wrongful φ -ing opens B up to moral complicity in that φ -ing?¹ I think that what it means for B to be morally complicit in A's

¹ Why say "opens B up to moral complicity" instead of just "B is morally complicit"? B can only be morally complicit in A's φ -ing if A goes on to φ . So B's deliberative duty failure opens her up to complicity but does not entail complicity, since A could choose not to φ . If A φ 's, then B will be complicit in that φ -ing.

wrongful φ -ing is for B to bear some moral responsibility for A's wrongful φ -ing in the absence of excuse or justification. This means that B bears some moral responsibility for the wrongful action (and the harm it constitutes or causes) in the world unless B's responsibility is excused or her deliberative duty failure justified. For example, when Brianna is morally complicit in Agatha's cheating, she bears some moral responsibility for Agatha's cheating. This does not mean that Agatha is not morally responsible for her cheating: in fact, I don't think that Brianna's being morally responsible for Agatha's cheating diminishes Agatha's responsibility for her own cheating in the least.² This means that Brianna is also *morally criticizable* for Agatha's cheating. And in cases in which A's wrongful φ -ing has victims, the complicit agent may owe apology or reparation to those victims, despite not committing the wrong against them herself. But when I call Brianna complicit, I do not mean that she is morally responsible or criticizable for Agatha's wrongdoing to the same degree as the Agatha is herself. Rather, the claim is that she is partly responsible *for Agatha's wrongdoing and its consequences*, even though usually we are not responsible for the actions or behavior of others.

The idea that failure in our deliberative duties can result in our being morally responsible for others' actions (and their consequences) is a *prima facie* puzzling claim. It does not often seem that we are responsible for *others'* wrongdoing, and it can be natural to think that my moral responsibility begins and ends with my own actions. Once someone else is acting, the consequences are on *them* and not on me. And even a reader who agrees with me that there are deliberative duties that govern our interactions with others and their practical deliberation might be tempted to deny this stronger and more controversial claim in favor of a different picture. For example, we might think that failure in one's deliberative duties results in moral responsibility for *that* failure but not for a further

² There are things that Brianna might do such that Brianna would be morally responsible and Agatha would not, or would have diminished moral responsibility. For example, Brianna might coerce or manipulate Agatha into cheating. But Brianna's having moral responsibility for Agatha's cheating does not on its own entail Agatha's having reduced moral responsibility for her cheating.

wrongdoing committed by another person. If this were right, Brianna would be morally responsible for her failure to talk responsibly with Agatha about what to do but *not* morally responsible for Agatha's cheating. I will raise and discuss this difficulty in Chapter 4, after developing more of the detail of the deliberative conception of complicity in earlier chapters.

In Chapter 1, I argue that the deliberative view can explain cases that other philosophical accounts of complicity cannot, including silence as complicity, complicity without difference-making, and complicity without intentional participation. For example, the deliberative view can explain why Brianna is complicit in Agatha's cheating, if Agatha goes on to cheat in the case described above. Brianna has failed in the deliberative duty she has to talk with Agatha about the reasons cheating is wrong—about why Agatha shouldn't cheat. So she fails in a deliberative duty she has with respect to Agatha's cheating, and this failure grounds her moral complicity in Agatha's cheating. This is so even though Brianna did not make a difference to Agatha's cheating nor did she intend that Agatha cheat.

The deliberative view is compatible with other philosophical accounts of moral complicity, provided that they are offered as sufficient and not necessary conditions on the phenomenon. That is, it is possible that some cases of moral complicity do not involve deliberative duty failure and instead involve some different grounding feature, such as intentional participation in the wrong action or making a causal contribution to the wrong action. These sorts of accounts will be explored in detail in Chapter 1, when I argue that neither intentional participation nor causal contribution is necessary for moral complicity, but a brief gloss may be helpful here. Accounts of moral complicity based on intentional participation typically require that the complicit agent intend to participate in the wrongful action in which they are complicit. Accounts of moral complicity based on causal contribution typically require that the complicit agent contributed causally to the wrongful action in which they are complicit, perhaps by increasing the likelihood that another agent would perform that

action. The account I offer here leaves open the possibility that some cases of complicity are grounded in one of these other features rather than in deliberative duty failure, and so it is offered in a pluralistic spirit. It could be that the best account of moral complicity is a disjunctive one that incorporates several different grounding bases for complicity.

But this does not mean that these bases of complicity are on equal footing. In Chapter 2, I argue for a default entitlement that agents have to act trustingly toward one another, thus focusing their attention on fulfilling their deliberative obligations rather than on ensuring that they do not do or say (or fail to do or say) anything that might causally contribute to wrongdoing. For example, imagine a case slightly different from our original. Not only will Agatha cheat even if Brianna discusses with her the reasons cheating is wrong, but in fact Brianna's talking with Agatha about why she shouldn't cheat actually *increases* the likelihood of Agatha's cheating. Brianna's default entitlement to act trustingly toward Agatha means that she is entitled to (and should) talk with Agatha about the reasons not to cheat anyway, despite this fact that such discussion will increase the likelihood of Agatha's cheating. This entitlement is in place as a default, so Brianna needn't first look into whether or by how much her discussion will increase the likelihood of Agatha's cheating. Rather, she is entitled to act trustingly toward Agatha unless the circumstances are exceptional and thus the entitlement is not in place.

I said above that the deliberative conception of complicity is compatible with other conceptions of complicity. That is, I offer it as a sufficient but non-necessary condition on moral complicity. While this is true, it is also important to note that deliberative complicity can arise in all sorts of cases, including those that at first might seem best explained by one of the other philosophical accounts of moral complicity. For example, the deliberative view can also explain cases of moral complicity that do not involve speech. Consider a case often thought to be a core case of complicity: Betsy is the getaway driver at a bank robbery; Adam robs the bank. Betsy

intuitively seems complicit in Adam's bank-robbing, though this case can seem difficult for the deliberative view. Even if Betsy had told Adam that robbing a bank is wrongful, it still seems that it is *driving* Adam to the bank that makes Betsy complicit in Adam's bank-robbing. To help explain this sort of case, in Chapter 3 I defend a theory of action according to which our conduct expresses our stance on moral issues. According to this theory, when Betsy drives Adam to the bank Betsy *acts-as-though* Adam's bank-robbing is permissible and so *expresses* that Adam's bank-robbing is permissible. In so doing, she violates her deliberative duty with respect to Adam's bank-robbing and so opens herself up to complicity in the robbery. More generally, an agent B's *acting-as-though* another agent A's wrongful φ -ing is morally permissible will sometimes constitute a failure in B's deliberative duty. And when it does, B is open to complicity in A's φ -ing.

Finally, Chapter 4 develops some of the details of the deliberative account of moral complicity, including exploring the connection between the explanation of complicity it offers and the scope of complicity that explanation implies. It also considers some difficulties and objections that arise for the deliberative view, including how the content of the wrongful action might make a difference to the scope of complicity and how we should think about the temporal element of complicity. For example, if I violate a deliberative duty with respect to A's wrongful φ -ing, it can seem implausible that I am complicit every time A φ 's, but it is difficult to see how the deliberative view can make sense of this intuition and how we should understand the connection between time and complicity more generally.

A deliberative account of moral complicity can help us to reimagine our understanding of our moral impact on the world and so help us make progress on difficult questions of our complicity in an increasingly global society. By helping others to see the moral truth, I can not only avoid complicity in wrongdoing but also support others who aiming act rightly. Through my speech and action, I express my stance on moral issues, and ensuring that those stances are both correct and in

line with my actual thinking is important not only to my own moral flourishing but also to the moral success of those around me.

Chapter 1: The Need for a Deliberative Conception of Moral Complicity

1.1 Introduction

Suppose that Albert solicits your position on promise-breaking, and you tell him there is nothing at all wrong with breaking a promise.¹ The next day, he wrongfully breaks a promise he had made to Cindy, and she accuses you of being complicit in Albert's breach. You might defend yourself against this charge by arguing that you didn't intend for Albert to break his promise: you didn't aim at or hope that Albert would act the way he did. Cindy is unsatisfied with this defense and claims that intentions don't settle the question of complicity: speaking in favor of promise-breaking is sufficient.

So you try a different tack and contend that your action did not make a difference to what Albert did: he would have broken his promise to her regardless of what you said. She remains unsatisfied. Perhaps he would have, but she argues that this *also* doesn't settle the question of your complicity. As she sees it, your action supported his promise-breaking, and that is enough to make you complicit, notwithstanding any difference it made or didn't make to Albert's wrongful action.

The main argument of this chapter is that Cindy could be right. Moral complicity requires neither intentional participation nor difference-making: you could be complicit in the absence of both. Of course, she may be wrong about the particular case. The context and content of your conversation with Albert seems to matter. To determine whether or not you are complicit, we seem to need information about the situation beyond the facts that you didn't intend nor make a difference to his promise-breaking. If more detail about the case matters (as it intuitively seems to), then views that rely only on difference-making or intentional participation are not capturing all of the facts relevant to moral complicity. We need more information about the situation and the

¹ You might know that promise-breaking is wrong or not; I think it we are obligated to give correct advice, not just advice we believe.

agent's action. Prevailing philosophical accounts of moral complicity cannot make sense of why this should be so.²

In this chapter, I will defend the *deliberative view* of moral complicity, an account that can incorporate—in fact, requires—more detail about what the potentially complicit agent did. While other accounts focus on the way in which a complicit agent is related to the moral failure of *another*—as intending or making a difference—this account focuses on a failure on the part of the complicit agent herself. In particular, it focuses on a complicit agent's failure in her deliberative duties, which are the obligations she has with respect to the practical deliberations of other agents. When we fail in our deliberative duties, we can become complicit in wrongful actions that those other agents go on to perform. By focusing on a failure of the complicit agent, the deliberative view of complicity can explain why the details matter: they make a difference to whether or not she has failed in her deliberative duty.

This chapter will proceed as follows. In §§1.2–1.3, I will show that accounts based in difference-making and intentional participation are extensionally inadequate. Then, in §1.4 I will propose an explanation for that inadequacy: such views try to determine complicity by considering only some sort of nonmoral relation between the potentially complicit agent and the relevant immoral action. But this strategy is, I think, impoverished from the start. A better account of moral complicity will look to the duties the potentially complicit agent *herself* is under, and §1.5 will sketch some of the details of one such account—the deliberative view of moral complicity. This account offers the resources we need to incorporate more of the morally relevant information into our philosophical understanding of moral complicity. §1.6 takes up a structural concern: why should we

² For example, see Kutz (2004) and Lepora and Goodin (2015). Two notable exceptions include Stilz (2011) and Frowe and Parry (2019). Stilz's discussion and Frowe and Parry's discussion help to explain why some of this detail matters, but their discussions are (by design, given their aims) limited to complicity only in certain contexts, not moral complicity more generally. I aim to defend an account of moral complicity that can explain why such detail matters to moral complicity generally speaking.

think that duty failure opens us up to complicity? In sum, I suggest that the deliberative conception of complicity is more satisfying than views based in difference-making or intentional participation. This leaves open the possibility that some cases of moral complicity are causal and not deliberative. That is, my argument here—that neither intentional participation nor difference-making is *necessary* to complicity—is consistent with the possibility that other forms of complicity are also possible. I will conclude that at least *some* cases of moral complicity are due to deliberative duty failure.

1.2 Difference-Making: Neither Sufficient nor Necessary

After a gunman with an assault-style rifle shot 48 people at a Walmart in El Paso, Texas in 2019, the superstore chain announced that it would discontinue the sales of short-barrel rifle ammunition, handgun ammunition, and handguns (McMillon 2019). This shift was significant, since Walmart made up a substantial portion of the ammunition market for decades (Hawkins and Krakow 2019). Walmart had been accused of complicity in prior mass shootings because guns purchased at their stores had been used to commit some of those atrocities (Honan 2019). In response, Walmart might reply to this kind of accusation that they are not complicit, since if they did not sell the guns and ammunition, the gunmen would simply have procured them elsewhere. Such a defense seems to rest on a difference-making view of complicity:

Difference-Making View:

In order to be morally complicit in some wrongful action φ performed by some agent A, an agent B has to perform some action ψ that makes a difference³ to whether or not A φ 's.

But a defense of this kind on Walmart's behalf is quite unsatisfying, and I think this is because moral complicity does not require difference-making. Even if it is true that the gunman would have purchased the weapons elsewhere, it seems to matter that he *in fact* purchased them from Walmart.

³ By “makes a difference,” I just mean that B's ψ -ing made A's φ -ing more likely. For a defense of a view of this sort in the case of legal complicity, see May (2010, 136).

But if this is right—if it *does* matter where he in fact procured the weapons, even in the face of other options being readily available—then difference-making is not a necessary condition on moral complicity.

The idea that difference-making is not a necessary condition on moral complicity seems even clearer in the case of an individual actor, as opposed to a corporate one. For example, consider a case adapted from one due to Bernard Williams:

George is offered help by “an older chemist”—let’s call him James—in securing a decently paid job in a laboratory that pursues research into chemical and biological warfare (CBW). George has a Ph.D. in chemistry, suffers from health problems, has young children to support, and has been having trouble finding a job. George refuses James’ help despite these difficulties because he is opposed to CBW. James confesses that part of his reason for offering to help George secure the job is because he is concerned about the person—let’s call him Bernard—who will take the job if George does not and is not inhibited by scruples concerning CBW. James believes that Bernard will push the research along with a greater zeal than George would (1973, 97–98).

It is a feature of the case that George will not make a difference to whether or not⁴ the research into CBW is conducted.⁵ And yet George doesn’t want to take the job: he feels that taking the job and working at the lab would make him complicit in the research, development, and deployment of CBW. Importantly, George’s feelings seem to make sense here. They don’t seem misplaced or mistaken but rather quite compelling. Intuitively, it *does* seem that working at the lab would make

⁴ He will make a difference, James predicts, to how quickly the research proceeds. But even if he refuses the job, the research will take place. Just as Walmart contends that if they don’t sell the guns and ammunition, someone else will.

⁵ This feature is purposeful: Williams uses the case to argue that act-utilitarianism misses something important about personal integrity and the role it plays in our understanding of moral theory. Williams’ point isn’t that act-utilitarianism gets the wrong answer about the case. That is, Williams thinks that we may come to agree that George ought to take the job (as act-utilitarianism would say he should). But he argues that act-utilitarianism gets this answer entirely too easily and too quickly and that this shows that the theory does not have space for the importance of personal integrity nor for the related idea that it matters to each of us whether evil enters the world “through me.” I agree with Williams that act-utilitarianism gets the answer entirely too quickly and too easily, but I am skeptical that the answer it gets is even correct: I do not think that George should take the job. However, on some readings of the Williams passage, what George does is not only move the research along more slowly than the person I have called Bernard would have but actually actively and purposefully slows the research. That is, he thinks of himself as undermining and sabotaging the research into CBW. If this is the right reading, it would make much more sense to me why Williams thinks that perhaps George could and should take the job, despite his serious opposition to CBW. In the argument in main text, I assume that though George would slow the research, he *does* conceive of himself as contributing to it and does not think of himself as a saboteur. Thus, the example in the main text may be different in morally significant ways from Williams’ George.

George complicit. This is further evidence that difference-making is not necessary for moral complicity.

Causation might still have a role to play in our theory: it may be that some kind of causal *contribution* is a necessary condition on moral complicity.⁶ For example, causal contribution might help to explain why it matters that the gunman *in fact* purchased the weapons from Walmart, even if it remains true he *could have* gotten them elsewhere. Walmart was a part of the causal chain that culminated in the wrongful killing of innocent people, despite not making a difference to whether that killing occurred. But causal contribution (even given the best account of causation) is certainly not a *sufficient* condition on moral complicity.⁷ I might drive by you on the freeway, distracting you because I happen to be driving your favorite car, resulting in your getting into an accident. Though I would be among the physical causes of your accident, I wouldn't be complicit. Causal contribution might be necessary for moral complicity, but it is certainly insufficient. What might be needed, instead of or in addition to casual contribution, in order to ground the moral complicity of agents?

1.3 Intentional Participation: Not Necessary

A view based on the potentially complicit agent's intentions might seem promising. After all, if George accepts the job, he intentionally participates in the production of CBW, despite his moral

⁶ Something like causal contribution seems to be what Gardner thinks is necessary for moral complicity (2004; 2007). Kutz, on the other hand, argues that casual contribution is *not* necessary for moral complicity (2000; 2007). As I will argue in Chapter 2, the more difficult problem facing a causal view of moral complicity is the guidance it provides a deliberating moral agent: she must think of all of herself as potentially morally responsible for all of the possible causal consequences of her actions, including those that go by way of the free actions of other agents, and this undermines the possibility of conceiving of those actors as genuinely free agents. But here I do not take myself to have shown that no version of causation could be a necessary condition on moral complicity, just that causal *difference* is not one.

⁷ The set of events we contribute to causally is overwhelmingly large, and we are surely not complicit in all of them. This idea is a major focus of the third season of NBC's *The Good Place*. See especially episodes 9, 10, and 11. The characters discover that *no one* has gotten into The Good Place (the show's approximate equivalent of heaven) in 521 years, due to the interconnectedness of the world and unintended consequences of even very small actions, such as buying a tomato. In future work, I hope to argue that deliberative complicity can help us to understand our moral impact in the modern world due to our involvement in a global economy.

opposition to its use.⁸ And it is precisely his intentional participation that seems to ground and make fitting his feelings that working there would make him complicit.

Intentional Participation View:

In order to be morally complicit in some wrongful action φ performed by some agent A, an agent B has to participate intentionally in A's φ -ing.

A view like this is well-suited to explaining complicity in cases of intentional participation without difference-making. George is reluctant to take the job *precisely because* it requires participating in the project of research into CBW, even though his taking the job will not make a difference to whether or not the research occurs.⁹

Importantly, the intentional participation view doesn't allow an agent to avoid complicity simply by changing how she conceives of what she is doing. As long as she is doing *something* intentionally, is aware that what she does contributes to the relevant wrong, and what she does *in fact* *does* contribute to the relevant wrong, she is complicit. For example, George might take the job and conceive of himself only as doing his part in whatever the company does. However, because he does his part intentionally, his part contributes to the company's aims, and he is aware of the company's aims and his contribution, the intentional participation view holds that he is complicit in the research, development, and deployment of CBW, even if participating in such research, development, and deployment is not an aim he would describe himself as having.

But views based on intentional participation encounter difficulty explaining complicity in cases in which agents seem complicit despite not participating intentionally. Though intention is a

⁸ It seems that participating in the production of CBW is something that George does intentionally but not something George intends, given his opposition to its use. (He would be happy to work for the company if it produced something else, for example.) For the distinction between what an agent intends to do and what she does intentionally, see Bratman (1987). Throughout, I will use "participates intentionally" to mean that the agent B participates in the relevant project X intentionally, as opposed to "intends to participate" for when the project X is what B intends to participate in.

⁹ To take a different example, the individual pilots involved in the bombing of Dresden, Germany during World War II were complicit in the atrocities that followed, even though no individual pilot's dropping a bomb made a difference to whether the firestorm took place (Kutz 2000, 122).

notoriously difficult philosophical concept,¹⁰ the debate surrounding it is immaterial to our current focus. My objection to views based on intentional participation does not depend on understanding intention in a particular way. Return to George's case, but suppose the facts are somewhat different. George takes the job and works at the company, but he doesn't know that it makes CBW. It's relatively easy to find out that the company makes CBW, and if George wanted to know he could, but he is content "doing his part in whatever the company does" and is unaware that it researches and produces CBW. In such a case, it does not seem that we can say that George participates *intentionally* in the research and production of CBW, since he is unaware that it is a project in which his company is engaged.¹¹ But George does seem complicit, especially if he could have and even *should* have known about his company's activities. If this is right, then George needn't participate intentionally in order to be complicit, and a view based on intentional participation seems unable to explain his complicity.

How might a proponent of an intentional participation view respond? Christopher Kutz, who defends a detailed and compelling view of moral complicity based on intentional participation, replies to a worry of this kind. According to Kutz, *everyone* who works at George's company is complicit in the deployment of CBW because the production of CBW is part of the *collective* end of the company, and that information is readily available. "So long as the decision to work with the company is voluntary, and information about the company's activities is available, *every* employee bears an accountable relation to the victims of CBW" (Kutz 2000, 157, emphasis mine). Kutz's solution shifts our focus from the aims of the individual agents to the collective ends of the company and holds the participating agents complicit in those collective ends, so long as information about those ends is readily available and employment at the company with those ends is

¹⁰ Cf. Anscombe (2000), Davidson (1963), Bratman (1987), among others.

¹¹ It is hard to see how an agent can intentionally participate in X if she is unaware of X. See Bratman (1987, 123).

voluntary. This formulation offers an important advantage: it can capture a lot of cases of moral complicity that involve culpable ignorance by those who are in positions to know about wrongdoing but remain unaware, perhaps hoping to secure plausible deniability.

This solution relies on the company's collective end and holds all employees at the company complicit. This means that it has difficulty explaining cases of complicity in wrongs that are not part of the collective *end* of a group and cases in which we might think that *some* but not all employees are complicit. For example, suppose that George's company has a terrible record and culture of sexual harassment and gender discrimination. This is a fact about it as a place of employment, but is not one of the collective goals of the company in any sense. Women are regularly passed over for promotions for which they are qualified, and jokes and comments of a sexual nature or disparaging toward women are a common occurrence. Many women experience the company as a hostile work environment. Who is responsible for this wrongful treatment of the female employees, and how should we understand the complicity of other agents at the company? Those responsible as *primary* wrongdoers tell jokes, make comments, or discriminate on the basis of gender. Of the remaining employees, it intuitively seems that some, but not all, will be complicit.¹² They may be complicit if they always remain silent or laugh when such jokes or comments are made, or if they notice (or should have noticed) the discriminatory hiring practices but failed to do anything to report or change them. Others may not be complicit, despite the ongoing harassment, if they do report the problematic culture, or if they make efforts to redirect the conversation away from such jokes or comments. An account of complicity needs the resources to distinguish between these two groups.

¹² Following legal discussions of criminal complicity we might refer to complicit agents as *secondary* wrongdoers, in contrast with the primary wrongdoers who commit the primary wrong (the wrong in which the secondary wrongdoer is complicit). Discussions of criminal complicity also sometimes call the primary wrongdoer the *principal* and the secondary wrongdoer the *accomplice*. See e.g. Gardner (2007), among others. I will reserve "secondary agent" or "secondary wrongdoer" for an agent who *is* complicit and use "potentially complicit agent" for an agent who opens herself up to complicity or an agent whose complicity is being considered/debated.

A case like this poses two problems for a view of moral complicity based on intentional participation. First, such a view doesn't seem to count *any* of the employees at the company as complicit, since sexual harassment and discrimination, though prevalent at the company, are not among the company's *ends*. We might try to address this first problem by expanding our notion of what the employees are complicit in beyond the company's ends or goals to *all* of the company's activities.¹³ As in the case of CBW, this would succeed in capturing the employees as complicit in sexual harassment. But adopting this solution runs into the second problem: it captures *all* of the employees as complicit, even though intuitively only *some* of them are complicit. A view based on intentional participation does not seem to have the resources to distinguish *amongst* employees of the company with respect to complicity other than (1) whether their employment is voluntary and (2) whether they knew or had reason to know of the company's activities. This leaves it unsatisfying as an account of complicity.

Before turning away from intentional participation views, it is worth considering one more way of spelling out the view that seems to address some of the problems we have considered. Suppose the view didn't require that B intentionally participate in A's φ -ing, since such intentional participation seems to require that B is aware of A's φ -ing. Instead, it might require that B intentionally do something and that A's φ -ing be among the reasonably foreseeable projects to which that something contributes.¹⁴ Note that in the revised George case, it is relatively easy to find out that the company makes CBW, and this seems important. If it was very difficult for even employees to determine that the company was involved with CBW, the case for George's complicity would be much less convincing.

¹³ With the same two caveats that (i) the individual employees know or should have known about those activities and (ii) their employment is voluntary.

¹⁴ By reasonably foreseeable projects, I have in mind the kinds of projects that B *should* be aware of, given the information available to him. For an example of a similar view concerning reasonably foreseeable consequences, see Scanlon (2008, 38–39).

Intentional Participation View (Reasonably Foreseeable)

In order for an agent B to be morally complicit in some wrongful action φ performed by some agent A:

- i. B has to intentionally ψ .
- ii. B's ψ -ing must contribute to A's φ -ing.
- iii. A's φ -ing must be among the reasonably foreseeable projects to which B's ψ -ing contributes.

This view seems on obviously better footing than the original formulation and seems well-equipped to handle our George cases, including both the original case and the modified version. Importantly, it seems well-positioned to handle cases that involve agents' willful and culpable ignorance about their contribution to wrongful actions, as it classifies agents as complicit in projects they contribute to if they could have reasonably foreseen their contribution to those actions and projects. *And it* seems to have the resources to distinguish amongst employees at the company in the sexual harassment case: those employees who could have reasonably foreseen that their activities contribute to sexual harassment and gender discrimination are complicit in it, but those who could not have (or whose activities do not contribute) are not complicit. So it will hold some but not all of the employees complicit in sexual harassment, which seems like a better result than on the views considered so far.

Unfortunately, although this modified intentional participation view does seem to have the resources to distinguish amongst employees, capturing some as complicit and some as not, it does not seem to draw this distinction in the right place. Some employees seem complicit on the basis of what they do *not* do, rather on the basis of what they do. For example, we noticed that some employees may be complicit if they always remain silent in the face of disparaging jokes or if they notice (or should have noticed) discriminatory hiring practices but failed to do anything to report or change them. Even though nothing they *do* contributes to the wrongful treatment, the intuition that they are complicit stands. And this really gets to the heart of the problem with views based on intentional participation: it just isn't the case that one needs to participate intentionally in order to be

complicit. This doesn't rule out the possibility of intentional participation as a *sufficient* condition on moral complicity: it may be that intentional participation is enough to make someone complicit in another's wrongful action. But because there are situations in which agents are complicit despite no kind of intentional participation on their part, it fails as a necessary condition on moral complicity.

1.4 Diagnosis: The Detail Matters

In §§1.2 and 1.3 I argued that views of moral complicity grounded in difference-making or intentional participation are extensionally inadequate because they lack the resources to incorporate some details about the cases that seem intuitively to matter to moral complicity. For example, when we consider an employee at George's company and try to determine their complicity, it seems that we want to know more than just whether they made a difference or participated intentionally. Did they laugh at the jokes? Did they stay silent? Did they make an effort to direct the conversation elsewhere? If these details about the situation matter (as they seem to), then we need an account of moral complicity that takes them into account.

What might such a view of complicity look like? First, notice that the views considered so far try to determine an agent's complicity by considering a nonmoral connection between the agent and the relevant morally valenced (wrongful) action. For example, the difference-making view tries to determine whether an agent is complicit by assessing whether she made a difference to whether the wrongful action occurred. Since difference-making is a nonmoral relation, this view is only able to incorporate as relevant to complicity information about the potentially complicit agent and her actions that is non-normative. But an approach of this kind is impoverished from the start because it fails to incorporate and accept as relevant input detail about the situation unrelated to difference-making. For example, suppose that even if George had spoken up rather than remained silent, the sexual harassment at his company would have remained unchanged. A difference-making view does

not capture his silence as relevant to complicity because it is not relevant to difference-making. Similarly, suppose that George's silence is not a case of intentional participation, as seems certainly possible in some cases. Then an intentional participation view will not capture his silence as relevant to complicity because it is not relevant to intentional participation. And yet his silence (and other details not relevant to difference-making or intentional participation, such as whether they laughed, redirected the conversation, or reported the harassment) seems relevant to the question of his complicity. Since such detail is sometimes irrelevant to difference-making or intentional participation, these views cannot capture them as relevant to complicity. But such detail *is* relevant to the obligations the potentially complicit agent is under, and an account grounded in these obligations will be able to incorporate such detail and make sense of why it seems important.

In the next section, I will introduce one such account—the *deliberative view* of moral complicity. This account claims that agents are complicit when they fail in a particular set of duties they have with respect to the practical deliberations of others.

1.5 A Deliberative Account of Moral Complicity

Recall that deliberative moral complicity arises when an agent B fails in a deliberative duty she has with respect to A's wrongful φ -ing and A goes on to φ . If I am right that this is a relevant category of moral complicity, then deliberative duty failure (when combined with A's performance of the relevant wrongful action) is sufficient for moral complicity. Here is a simplified statement of the view:

Deliberative View of Moral Complicity:

B's failure in her deliberative duty regarding A's φ -ing is *sufficient for* and *explains* B's moral complicity in A's wrongful φ -ing, if A φ 's.

Suppose George tells his coworker that he thinks their workplace culture is fine, and that coworker goes on to make a comment disparaging toward women at the next staff meeting. The deliberative

view of moral complicity holds that George is morally complicit in his coworker's wrongful treatment of the women (and men) at the meeting. Why? Because in saying the workplace culture is fine when it is not, George failed in a deliberative obligation he has with respect to the wrongful treatment of women at the company.

In general, the deliberative view holds that an agent is complicit when she fails in a deliberative duty she has with respect to A 's φ -ing and A goes on to φ . What are our deliberative duties?¹⁵ Some are familiar, others less so. Examples include taking reasonable care¹⁶ to speak accurately and exercising due care with respect to the mental contents of others.¹⁷ *Sincere* speech matches what one believes, while *accurate* speech matches the way the world in fact is.¹⁸ Taking reasonable care to ensure accuracy includes having adequate evidence for what one says.¹⁹ The duty of care with respect to the mental contents of others requires that we take reasonable care to ensure that others' false beliefs are not formed nor confirmed on the basis of our speech, conduct, or omission.²⁰ And it sometimes requires that we speak up or make corrections to others' false beliefs.

¹⁵ I think of deliberative duties as moral duties that we have with respect to the epistemic domain. However, if Jennifer Lackey is right that epistemic duties can be interpersonal and not strictly doxastic, then it is possible and even likely that the deliberative duties have both a moral and epistemic component (2020). I take no stand on the debate concerning the distinction between moral duties governing the epistemic domain and "pure" epistemic duties here. For an argument that the epistemic duties are intrapersonal and strictly doxastic, see Wrenn (2007).

¹⁶ Cf. Grice's "Maxim of Quality": "Try to make your contribution one that is true." Of course, the maxims of conversation that Grice has in mind don't seem to be *moral obligations*, though they can be exploited by those who aim to deceive (1991, 27). Note that on my account *trying* to make one's contribution one that is true won't satisfy the demand; rather, in order to fulfill the duty an agent will have to meet a *reasonable person standard* of truthfulness. For further detail about the distinction between speaking sincerely and speaking accurately, see Williams (2002, 126).

¹⁷ For a compelling and insightful argument regarding the requirement of sincere speech and the wrong of the lie see Shiffrin (2014, Chapter 1).

¹⁸ So speaking *inaccurately* by accident is possible in a way that speaking insincerely by accident is not, since one could be mistaken about the facts—about the way the world is—in a way that one could not be mistaken about what one is aware one believes. As a matter of fact, speaking accurately is also possible by accident in a way that speaking sincerely is not, since one could actively intend to say what is untrue but be mistaken about the truth and so speak accurately. Such a case would *not* count as fulfilling the duty of due care with respect to the mental contents of others, since the person did not in fact take due care (cf. Herman (2019)). But this doesn't mean that their action would be impermissible (cf. Scanlon (2008)).

¹⁹ What counts as having adequate evidence will vary along familiar lines, including the context and content of the information presented. A physician's assertion to a patient that a certain procedure is safe will require a higher threshold of evidence than a bystander sharing the time or answering whether the bus has passed.

²⁰ As I understand this duty, it does not rule out the possibility of permissible deception in cases of e.g. non-lying deception intentionally deployed to protect one's reasonable sphere of privacy. See Shiffrin (2019, 79).

When am I required to speak up? Perhaps I've been asked for advice about what to do. Perhaps the context makes a particular content salient between me and another agent. I won't aim here to give a complete description of when exactly the duty requires that we speak up: rather, my aim is to give a sense of the shape of the duty so we can assess its role in the structure of moral complicity. In fact, substantive disagreements over whether a deliberative duty is in place in a particular case are possible and even expected. But these kinds of substantive disagreements are independent of my structural contention that deliberative duty failure opens one up to moral complicity.²¹ The important point here is that looking to the complicit agent's duty failure provides needed resources in considering the question of complicity and allows us to incorporate the detail that seems significant to complicity but was neglected by the other accounts.

Suppose George's coworker and friend asks George what he thinks of the culture around the office.²² By asking George what he thinks, his coworker makes the culture of the office salient between them. If George says that the culture is perfectly fine, he fails in his deliberative duty to take reasonable care to ensure that others' false beliefs are not formed nor confirmed because of his conduct. Note that this is so even if George *believes* that the culture is perfectly fine (and so speaks sincerely), since the duty requires meeting a reasonable standard of accuracy, and the culture is *not* perfectly fine (by assumption). The deliberative view holds that George's failure opens him to complicity if his coworker then makes some disparaging comments.

The deliberative view focuses on the obligations that the potentially complicit agent is under with respect to the practical deliberations of others. One upshot of the view is that she acts wrongly if she fails in her deliberative duty, even if the primary agent does not go on to act wrongly himself.

²¹ Scanlon makes a similar point about objections to his arguments defending contractualism. An objector might disagree with Scanlon's evaluation of what principles people could reasonably reject, but this would not be an objection to contractualism as such but only to Scanlon's evaluation of a particular case. Scanlon calls this kind of disagreement a "substantive disagreement," and I follow his use of the phrase (2000, 186, 205, 238).

²² It is possible that the power dynamic between George and his coworker affects things between them, including George's deliberative duties with respect to his coworker's potential actions. See "excusing conditions" below.

She will not be *complicit* in such a case, since there won't be an action to be complicit in.

Let us consider a few more difficult cases and how the deliberative view might handle them. First, there is a question of how to handle cases of *excusing* conditions with respect to the deliberative duties. Perhaps the possibility of A's φ -ing has been made salient between A and B but the power dynamics between them leave B uncomfortable or vulnerable if she explains to A why it would be wrong for A to φ . In such a case, we might say that B is excused from her deliberative duty. While we might be tempted to say that B is not complicit in such a case, I think we should instead say that her complicity is also excused. This better captures what makes cases of marginalization and oppression distinctively wrong: marginalized and oppressed persons often feel complicit in their own marginalization and oppression.²³ Though their complicity surely is often or even always excused, there are good reasons grounded in the reported experience of the marginalized and oppressed to want to hold onto the idea that they are complicit.

Second, how should a deliberative account handle cases of *justification*? Perhaps B seems to have a deliberative duty with respect to A's φ -ing because A's φ -ing has been made salient between them, but because of features of the context B in fact does not have such a duty. For example, B might not need to fulfill her deliberative duty at a particular time and place despite the content's salience because A and B have a background history of having discussed the issue, and B has fulfilled her deliberative duty during these prior discussions. Raising the issue again in the moment may be condescending or browbeating. It may even be that sometimes B *should not* engage with A regarding A's potential φ -ing, such as if B is in a place of authority over A, and giving advice (however well-intentioned) to A would approach or even constitute coercion.²⁴

²³ For example, see (among others), Dotson (2018, 196).

²⁴ For example, if B is A's boss, B's advice-giving may come across as a contingency announcement even if it isn't intended to be: "Keep your promises or I will fire you!" A contingency announcement is an announcement on the part

There are lingering worries here concerning the fair distribution of epistemic labor and moral demands in nonideal theory. Full consideration of these interesting questions is outside the scope of my discussion here, though I hope to engage with them in future work. I suspect that cases of epistemic exploitation will often constitute cases of excuse or justification.²⁵ And epistemic moral demands in nonideal theory seem to raise their own interesting and difficult questions, distinct from non-epistemic cases such as e.g. beneficence.²⁶ Thinking about such demands in the epistemic domain may even help to illuminate the non-epistemic cases. For now I set these issues to the side and turn to my argument concerning the structural contention of the deliberative view of moral complicity—namely that failing in one’s deliberative duty opens one up to complicity.

1.6 The Connection

The structure of the deliberative view of moral complicity can seem surprising: why would a failure in a deliberative duty open one up to complicity? After all, duty failure isn’t generally transitive: if I break a promise to you and you break a promise to C, it isn’t as though I break a promise to C.

My main structural suggestion is that failure in one’s deliberative duty opens one up to imputability for consequences that follow from that failure. Importantly, it is part of this idea that

of one party (B) that what she will do is contingent on what the other party (A) does. Such announcements are often coercive and impermissible. See Pallikkathayil (2011).

²⁵ Epistemic exploitation occurs when persons of privilege expect or demand that members of marginalized groups “produce an explanation about the nature of the oppression they face” (Berenstain 2016, 570). For example, a man might demand that a woman explain her experience of gender oppression in the workplace. Because the deliberative duties I have defended require that we understand ourselves as responsible for helping others come to see the moral truth, it can seem as though the deliberative view of moral complicity implies that marginalized groups are required to submit to epistemic exploitation in order to avoid complicity. (So in the example, the woman has a deliberative duty to provide the explanation.) I think the view is not committed to this result but do not offer a full defense against this worry here.

²⁶ The overdemandingness problem facing the duty of beneficence is well-recognized yet notoriously difficult to articulate in a precise way. See Murphy (2003) for a discussion of the difficulty as well as a proposal regarding what we ought to do when so many around us fail to do as they ought. In future work, I hope to investigate these questions as they bear on the epistemic case, as I think they will help to illuminate how we should respond to issues such as Fake News and the spread of misinformation, including on social media, as well as our responsibilities and complicity surrounding such issues.

those consequences would *not* be imputable to you if you had fulfilled your duty. This idea is not a new one: Kant expresses a version of it in his (in)famous “On the supposed right to lie from philanthropy” in which he argues that lying is wrong, even to a murderer at one’s door.²⁷ The passage is variably criticized as Kant’s overcommitment to rigorism or defended as not really entailed by Kant’s moral philosophy, but for my purposes we need not settle this debate.²⁸ What is crucial is Kant’s surprising claim that the person who lies to the murderer at the door can be “prosecuted as the author” of the victim’s death (1999, 612). A wrongful lie opens up the liar to imputability for consequences that it seems absurd to assign to her. She is the *author* of the victim’s death, even though she did not kill the victim and even *hoped* that he would *not* be killed. After all, her reason for lying in the first place is to prevent the would-be murderer from finding his intended victim. Kant’s suggestion here is that the agent’s duty failure (wrongful lying) opens her up to imputability for consequences that would not be imputable to her if she had fulfilled her duty instead,²⁹ including consequences (such as murder) that are a result of the free actions of other agents.³⁰ When we follow our duty, the negative consequences that follow are not imputable to us, but when we *fail* in our duty, those consequences *are* imputable to us.

A worry for Kant’s view is that it is too broad: it implies that we are morally complicit in *any* wrongful action that causally follows from a failure in any of our duties. My suggestion is much narrower in scope: failure in deliberative duties opens us up to responsibility for a *particular* set of consequences—the performance of a wrongful action by another person—when those

²⁷ See Kant (1999).

²⁸ For a compelling argument that Kant is *not* committed to its being wrong to lie to the murderer at the door see Korsgaard (1986). Shiffrin argues that it is not wrong to misrepresent to the murderer at the door (about some content) and goes further than Korsgaard, arguing that it is also not wrong to misrepresent to the murderer *next* door (2014, Chapter 1). For further illuminating discussion of this passage in Kant, see Herman (2011) and Wood (2008).

²⁹ Cf. Herman (2020).

³⁰ For a contrasting view, see Ruth Bader Ginsburg’s dissent in *Burwell v. Hobby Lobby*, in which she argues that complicity is not possible when the causal chain is “interrupted by independent decision makers” (2014).

consequences are connected to that failure in a particular way.³¹ It is *because* George fails in his deliberative duty with respect to his coworker's disparaging comments that he is complicit in his coworker's subsequent disparagement. But this failure doesn't open George to complicity in e.g. his coworker's theft, even if George's comment happened to make a difference to the theft. Our deliberative duties pick out what we owe with respect to the deliberations of others, and so failure in them opens us to complicity when others act wrongly. But the deliberative duties have the resources to make more fine-grained distinctions than a simple asymmetrical view of consequences can make. Because they can and do incorporate more morally relevant information, they attach a narrower scope of complicity to agents for duty failure. A failure in one's deliberative duty doesn't open one up to complicity *no matter what* the other agent goes on to do. Rather, it opens one up to complicity only if the other agent goes on to act wrongly with respect to the content that corresponds to one's deliberative duty failure. In general, B opens herself up to complicity in A's φ -ing *when and because* she fails to fulfill her deliberative duty with respect to A's φ -ing.

How does this solve the puzzle with which we began? If I break a promise to you and you break a promise to C, it *still* isn't as though I break a promise to C. But what if I tell you that breaking a promise to C is perfectly fine to do, that there is nothing at all wrong with breaking a promise? I still won't have broken a promise to C, but it *does* seem like C has a complaint against me. And this is some evidence my failure in my deliberative duty with respect to your promise-breaking opens me up to complicity in that promise-breaking. But suppose I tell you that promise-breaking is perfectly fine, and you go on to steal from C. Now it seems as though C does not have a complaint against me and that I am not complicit in the theft. So the idea that duty failure can open

³¹ In this way, my account of moral complicity is broadly Kantian in spirit, but I make no claim that it is in fact a view that Kant would have or did hold. My aim is a consistent and compelling view of moral complicity, and it is inspired by Kant but does not claim to be an interpretation of Kant's text.

one up to complicity, though puzzling and of course not fully defended here, can help to capture some of the intuitively plausible landscape of moral complicity.

1.7 The Devil's in the Details

I have argued that views based on difference-making or intentional participation lack the resources to take into account information about cases that intuitively seems relevant to determining moral complicity. Both accounts seem to make a similar mistake: they rest complicity on a nonmoral connection between a potentially complicit agent and the morally valenced (wrong) action of the primary agent. A better account of moral complicity will consider the obligations violated by the potentially complicit agent. We are now in a position to see that the deliberative view is one such account. It rests complicity on facts about the deliberative obligations the potentially complicit agent was under. And since the details we noticed as intuitively relevant to complicity *are* relevant to those obligations, it effectively captures them as relevant to complicity.

And yet we might wonder: if this detail is so important to complicity, why did difference-making and intentional participation accounts seem promising and attractive in the first place? I think it is because each of them is closely connected with deliberative duty failure, and so each can seem as though they are on the right track. In many cases in which B does not make a difference to A's φ -ing, B also doesn't fail in her deliberative duty with respect to A's φ -ing. This is because if B does not make a difference to A's φ -ing, B was often not in a position to have or fulfill a deliberative duty with respect to A's φ -ing. This means that difference-making can seem like it is on the right track as an account of moral complicity, not because difference-making itself is grounding moral complicity but because it is pointing toward something that is—problematic deliberation contribution via deliberative duty failure.

Something similar can be said about intentional participation. It can seem promising as an account of moral complicity because if an agent is intentionally participating in the wrongful action of another agent, it is likely that she is violating her deliberative duty in doing so. Intentional participation is overly narrow as an account of moral complicity as a phenomenon, but it will often constitute problematic deliberative contribution since intentional participation in another's wrongful action will often communicate endorsement of that action. Really, intentional participation seems better suited to explaining cooperating in wrongdoing or group wrongdoing than to explaining moral complicity. But since we might expect that those cooperating in wrongdoing fail in their deliberative duties associated with that wrongdoing, it may very well be that intentional participation is *sufficient* (though not necessary) for complicity.

Though neither difference-making nor intentional participation fully capture moral complicity as a phenomenon, they can seem attractive because they are closely connected to what matters to moral complicity, namely what we owe with respect to the deliberations of others.

1.8 Conclusion

In this chapter, I have defended a deliberative view of complicity on the grounds that it is in a better position to capture and explain the intuitively plausible landscape of moral complicity than views based in either difference-making or intentional participation. Difference-making and intentional participation views are unable to count as significant to moral complicity details about the cases that intuitively seem relevant. One way to make sense of the intuitive importance of such detail is to notice that it matters to the obligations the potentially complicit agent is under independently of whether the primary agent goes on to act wrongly. This means that the deliberative account of moral complicity, since it focuses on such obligations, can incorporate this intuitively relevant information as significant to moral complicity.

The deliberative account holds agents complicit when they fail in their deliberative obligations with respect to a particular agent and a particular action. A consequence of this view is that the potentially complicit agent acts wrongly when she fails in her deliberative duty *independently* of whether the primary agent goes on to act wrongly. Of course, she won't be *complicit* unless the primary agent acts wrongly, since there won't be a primary action in which to be complicit. But she still will have failed in a deliberative duty she had, and so she will have acted wrongly. She has gotten "lucky" with respect to moral complicity, but this does not mean she hasn't acted wrongly.³²

A worry for the deliberative view of complicity is that it is yet another sufficient condition on moral complicity rather than a necessary one. If so, it is incomplete as an account of moral complicity. This possibility is consistent with the argument of this chapter. I have argued that difference-making and intentional participation are each unnecessary and that violating a deliberative duty can open an agent to complicity and so is sufficient to ground moral complicity in the absence of either difference-making or intentional participation. But can the deliberative account possibly spell out a necessary condition? What about seemingly classic cases of complicity, such as driving the getaway car or providing the murder weapon? In Chapter 3, I argue that the deliberative account is more expansive than it might seem because our actions and silences often express our stances on significant moral issues.³³ And yet it is worth noting that even if this account cannot be made to work in the end, discovering a new sufficient condition is still interesting and important: the best view of moral complicity surely includes the story of our interpersonal deliberative obligations.

A deliberative account can help us to reimagine the moral landscape of complicity. While I may still reasonably worry about the difference I make to wrongdoing or injustice or about my

³² Compare Nagel's negligent truck driver, who fails to check the brakes but gets morally "lucky" in the sense of not hurting anyone despite his negligence (1979). The driver has still acted wrongly because negligently, but he isn't responsible for hurting someone since no one got hurt.

³³ See Anderson and Pildes (2000) and Helmreich (2015) for related views of action.

intentional participation in wrongful actions, to allow these to be the sole focus of my energy and attention is to miss a crucial aspect of my moral impact on the world. By helping others to see the moral truth, I can not only avoid my own complicity in wrongdoing but also support others who are aiming to act rightly. Ensuring that others' false beliefs are not formed nor confirmed on the basis of my speech or action is important not only to my own moral flourishing but also the moral success of those around me.

Chapter 2: Deliberative Duties and Our Default Entitlement to Adopt a Deliberative Stance Towards Others

2.1 Introduction

Most theories of moral complicity require that an agent who is morally complicit in the wrongful action of another agent causally contribute to that action.¹ For example, Bradley may be complicit in Agatha's manipulation of Charlie because Bradley taught Agatha how to manipulate someone. If Bradley hadn't taught Agatha how to manipulate, Agatha would not have manipulated Charlie, and so Bradley is complicit in her manipulation of Charlie in part due to his causal contribution to her manipulation. There is also general agreement that causal contribution, though necessary, is not sufficient for moral complicity. This is because causation is an overly permissive relation: Bradley might inadvertently bump Agatha, and the bump might suffice as a causal contribution to her manipulation of Charlie. But intuitively an inadvertent bump does not make Bradley complicit in that manipulation. Because causation is understood as necessary but insufficient for moral complicity, theorists then search for what might be required over and above causal contribution.

I think this focus on causation is mistaken. As I argued in Chapter 1, I think that some cases of moral complicity do not require causal contribution. If this is right, it raises a question of the relationship between different sorts of complicity. In this second chapter, I argue that moral complicity comes in (at least) two fundamentally different kinds: deliberative moral complicity and causal moral complicity. As we have seen, when an agent is morally complicit in the deliberative sense, her moral complicity in the wrongful action is grounded not in any causal difference she made or failed to make to the other agent's action but rather in a failure on her part to fulfill what I will call her deliberative duty. But this leaves open the possibility that an agent might become morally complicit in some cases due to her causal contribution, even if she fulfills her deliberative obligation.

¹ For example, see Gardner (2007), Sartorio (2007), and Petersson (2013). For a notable exception that has greatly influenced my own thinking on the topic see Kutz (2007; 2000).

And this raises a problem for the agent: should she focus her attention on fulfilling her deliberative duty and risk causal complicity or avoid deliberative engagement so as to ensure avoiding causal complicity? In this chapter, I defend a default entitlement to engage with others in what I call the deliberative mode. When B engages with A in the deliberative mode, B talks with A about what he should do and why without first looking into how she thinks A will act based on their conversation. If I am right that this default entitlement is often in place, then agents should focus on fulfilling their deliberative obligations. And they should do so because doing so is a way of respecting other agents as practical deliberators.

In §2.2, I develop an account of “practical co-deliberation,” a special sort of interaction between two agents in which each acknowledges that the other is a person who can act for the reasons she has. Engaging in practical co-deliberation with another agent is a way of respecting her *as* an agent. Such engagement is justified by a default entitlement to adopt a deliberative stance toward other agents, to treat them as agents who can act for the reasons they have rather than as things to be managed. Because practical co-deliberation is valuable for its own sake, it grounds some duties that govern interpersonal interaction around deliberation. Assuming that the default entitlement is in place, an agent who causally contributes to the wrongful action of another will not be morally complicit in that wrongful action if she fulfills her deliberative duty.²

2.2 The Value of “Practical Co-Deliberation”

There is a valuable way that two people can interact when that interaction is based on the mutually acknowledged premises that (i) both of them have things that they should do and (ii) they can help

² As we will see in Chapter 3, it can be a complicated matter whether a deliberative duty has in fact been fulfilled. What might seem like causal contribution could also constitute deliberative duty failure. For example, if I counsel you not to φ but then cover it up or make it easier for you to φ , I may still be complicit if the default entitlement is in place because these actions constitute a failure in my *deliberative* obligation. See Chapter 3 for more detail and an argument.

each other to think about what each of them should do.³ Talking with each other about what to do and why is a good way that two people can interact with one another in part because it is one way that they can acknowledge each other as persons. When B talks with A about A's reasons for or against φ -ing, B recognizes A's deliberative capacities. By mentioning to A the reasons B takes to bear on whether or not A should φ , B communicates to A that B takes A to be a reasoner who can think about what to do and why to do what's to be done. In turn, by taking B's mentioning of the reasons as an opportunity to look again into whether to φ and why or why not he should φ , A acknowledges that B is a reasoner who can help A to think about what he ought to do. In short, they communicate to one another that they take the other to be a reasoner who can think about what to do and why to do it. Let me name the sort of interaction in which two people discuss together what one of them should do and what her reasons are for doing it "practical co-deliberation."⁴

Practical co-deliberation is good for its own sake and not only good for the consequences that it may bring about. In particular, it isn't good just because when it happens it is likely to increase the probability that the individuals who interact in this way will then go on to act well. Rather, it is good because it constitutes agents respecting each other as autonomous reasoners. The agents do not just tell each other that they respect each other's deliberative capacities but also actually exercise that respect. In sincerely taking up together the practical question of what A should do, they engage as practical reasoners together.

³ We can recognize the value of such interaction without assuming that the agents couldn't have come to understand how they should act if they didn't have each other's help. The idea is that they *can* help one another come to know what they should do and that this kind of deliberative interaction is good.

⁴ I am using the phrase in a way that is different from how it is often used. The object of this kind of co-deliberation is still *individual action*. The two agents are not, in the cases I have in mind, co-deliberating in the sense of deciding what they together will do. Rather, they are talking with each other about what each should do and the reasons each has. A conversation between two agents would count as a case of practical co-deliberation, in my sense, even if the topic of the conversation were what *one* of the agents should do. In fact, these are my central cases.

Because practical co-deliberation is one way reasoners show each other respect as reasoners, it will sometimes be the appropriate way for agents to interact with one another. And when that's true, agents will have some associated deliberative duties. For example, agents should make this kind of interaction with one another possible by raising practical questions for mutual discussion. They should talk with each other about the reasons they have for doing what they should. Some such duties will have to be imperfect: it will not be a requirement that agents *always* raise practical questions for mutual discussion, nor should they always talk with each other about the reasons they have for doing what they should. And this is not just because it will not be possible always to be involved in practical discussions on pain of failing to get anything else done—though that's true. Rather, there are contexts and relationships within which certain practical questions are particularly salient, and contexts and relationships within which it isn't appropriate to discuss certain questions. We are not always obligated to intrude into each other's lives in this way, and sometimes we are obligated not to. But sometimes we *are* so obligated.

One familiar example of a deliberative duty is the norm of sincerity. In one sense, it is obvious that the norm of sincerity governs practical co-deliberation. After all, the norm of sincerity governs speech, and two agents who are practically co-deliberating are speaking with one another. Violating the norm of sincerity will be criticizable in a situation of practical co-deliberation for all of the reasons such a violation is normally criticizable.⁵ But it will also be criticizable on further grounds: violating the norm of sincerity interferes with the possibility of genuine practical co-deliberation. And so the value of practical co-deliberation is an additional reason not to violate the norm. Since genuine co-deliberation involves two agents talking with one another about what each should do, it requires that each says what she takes to be the case about what the other should do and about what reasons the other has. And so, if one of them says what she doesn't take to be so,

⁵ For example, see Shiffrin (2014, Chapter 1) for a compelling argument that lying is always wrong.

she won't be co-deliberating.⁶ And if the circumstances are such that co-deliberation is called-for, she will be directly criticizable for failing to co-deliberate and not just for violating the norm of sincerity.

When is it called-for? Certainly not always. I do not think we should go around telling just anyone what they should do and why, and a view that recommended we should would be a bad one for that reason. Rather, there are times when co-deliberation is appropriate and times when it is not. And this is not (only) a point about epistemic modesty. We might be justifiably certain about what someone ought to do in a situation and still have good reason to refrain from talking with them about the reasons they have to do as they ought. When should B stay out of it, and when should B talk with A about his reasons? I am skeptical that any set of necessary and sufficient conditions is possible, but we can articulate some of the relevant factors, including the relationship between A and B, the connection between that relationship and the considered action, and the context of the interaction.

Suppose that A and B are agents and that A is considering whether or not to φ . B knows that φ -ing is wrong.⁷ To be more precise, B knows that if A were to φ in the circumstances, A would act wrongly. What are some of the features that help to pick out a situation in which it is appropriate for A and B to engage in practical co-deliberation?

Their relationship will matter. And exactly how it will matter is a bit complicated. It isn't just that people who are closer with one another should be more ready to engage in practical co-

⁶ Again, the point here is fundamentally non-consequential. It may be that violating the norm of sincerity makes the success of practical co-deliberation less likely (though this seems contingent: it also may be that it makes it *more* likely, if success were measured in terms of the person simply performing the right action or avoiding the wrongful one), but that is not my criticism of such behavior here. Rather, my point is that violating the norm of sincerity makes genuine practical co-deliberation *impossible*. This idea is connected with Kant's contention that the person to whom one makes a lying promise "cannot possibly agree" to the way he is being treated (1999, 80). The claim is a logical one, not contingent nor dependent on good consequences. (Particular thanks to Barbara Herman for helping me understand this *Groundwork* example.)

⁷ I am assuming here that B is correct—were A to φ , A would act wrongly.

deliberation with one another. Rather, I think the sort of relationship that A and B have and the connection that relationship may have with φ will bear on whether it is appropriate for the two agents to engage in practical co-deliberation. For example, if A and B are co-workers and A is considering φ -ing, where φ is an action that is fundamentally work-related, B may be especially well-situated to engage in practical co-deliberation with A about whether to φ . And she may even have a special obligation to do so. But if φ is not work-related and is instead about the state of A's marriage, B may have an obligation not to interfere or offer advice or suggestions about what A should do.

Within the context of the relationship, A and B's previous interactions concerning φ will matter. If A and B have previously discussed the practical question under consideration, and if B knows that A already knows her stance on the matter and what she takes the reasons to be against φ , B may have good reason to refrain from taking up the matter with A yet again. Otherwise, arguing against φ may result in a sort of brow-beating.

It also matters whether A raises the question of whether he should φ as a question for B. If A raises the question or asks for help or advice regarding it and invites B into a conversation with him about what to do, it will sometimes be appropriate for B to accept that invitation and genuinely practically co-deliberate with A, even if A and B disagree about whether A's φ -ing would be wrong.

These considerations are meant to give a sense of when practical co-deliberation would be appropriate and a sense of some of the duties that arise from the value of practical co-deliberation. They do not give a full account of when practical co-deliberation is called-for, but I think some discretion on the part of agents is built into the value of the duty. This raises a further question: why is readiness to engage in practical co-deliberation a duty we have?⁸ One way to answer this is to

⁸ At this point, my reader may be skeptical that interfering in one another's lives in the way I suggest is *ever* appropriate. Someone who is a complete individualist about moral responsibility, I take it, will not be moved by my arguments here. But I do think there is a virtue to being involved in one another's lives an appropriate amount. For an intriguing

point to the value of the co-deliberation itself. If something is valuable, sometimes and under the right circumstances we should be willing to bring it about. I think this is at best only part of the story about why agents should be willing to engage in practical co-deliberation when it is called-for. The better answer returns to the value of practical co-deliberation. Recall that practical co-deliberation is valuable because it is a way of recognizing the personhood of another agent as an authoritative reasoner. By engaging in practical co-deliberation, two agents communicate to each other that they take each other to be persons who can act for the reasons they have.

So where do we stand? I have argued that practical co-deliberation is a valuable way that two agents can interact, and that it includes discussing together what one of the agents should do and why. I have suggested that it is a bit complicated to figure out when practical co-deliberation is appropriate or called-for. Further, I have argued that when practical co-deliberation *is* called-for, its possibility and aptness generate duties for those agents. First, they have, as always, a duty of sincerity that governs their speech. Second, they have a duty to be prepared to talk with each other about what do to and why: they have a duty to engage in practical co-deliberation. In the next section, I will argue that these deliberative duties and the value of practical co-deliberation offer informative lessons to our thinking about the deliberative conception of moral complicity.

2.3 Deliberative Complicity

In general, when B is practically co-deliberating with A, B doesn't predict either that A will act well or that A will act poorly. She makes no prediction about how A will act, understanding herself and

discussion of manipulateness and the virtue that may be associated with it that has greatly influenced my own thinking in this area, see Baron (2003) From page 48: "By contrast, the person who has the virtue corresponding to manipulateness—a virtue for which we do not, I believe, have a name—knows when it is appropriate to try to bring about a change in another's conduct and does this for the right reasons, for the right ends, and only in instances where it is warranted (and worth the risks), and only using acceptable means."

her own speech not as causally contributing to A's either φ -ing or not φ -ing but instead as collaborating in a discussion that contributes to A's deciding what to do.⁹ And because she understands herself this way, as a rational contributor to A's deliberation about what to do, she discusses with A reasons not to φ . Since B knows that φ -ing is wrong, she tries to persuade A not to φ . But she doesn't try to persuade A not to φ by mentioning just anything that might convince A not to φ . Rather, she tries to persuade A not to φ by mentioning precisely the reasons that A has not to φ . In so doing, perhaps she helps A to come to know his reasons not to φ and even to decide not to φ . But perhaps she does not. Suppose we call B acting toward A in the way described here as B "adopting a deliberative stance toward A" or "engaging with A in the deliberative mode". Alternatively, when B instead aims at getting A to act in a certain way, by understanding herself as a causal contributor to A's action, B adopts a "causal stance toward A" or "engages with A in the causal mode." This will often involve predicting how A will act conditional on what B does and then B's tailoring her actions accordingly.

When B is in the deliberative mode, B trusts A to take their conversation seriously.¹⁰ But this does not mean that B believes that A will hear the reasons she presents and act based on those reasons. She *may*, of course, believe that A will act on his reasons, once they discuss them. But getting A to act in a certain way is not B's only nor even her primary motivation for talking with A about the reasons he has.¹¹ After all, if it were, it might make sense to employ other means to get

⁹ It is worth being precise here: B may have a belief about how A will act conditional on what B does. And B's having such a belief won't preclude B's acting toward A in the way that I have in mind. What is essential is that B's prediction about how A will act conditional on what B does isn't part of the explanation for why B acts as she does. She adopts a deliberative stance towards A because of her entitlement to do so, because doing so is one way to treat A as a person who can act for reasons, and not because she predicts that so acting is likely to result in the best outcome.

¹⁰ Importantly, I do not think that co-deliberation will have to involve B's trusting A not to φ . This is because B's trusting A not to φ may require that B *believe* that A will not φ . (For a compelling argument that genuine trust requires belief, see (2008).) B's engaging with A in the deliberative mode, in the way I have in mind, is compatible with B's believing that A will φ . What is essential is that the question under discussion is whether A should φ , and both parties are seriously involved in a mutual discussion about that question, despite their beliefs about what A will or will not do.

¹¹ For simplicity and consistency, I refer to A using male pronouns and B using female pronouns throughout this second chapter.

him to act that way, perhaps means more likely to be successful at achieving that aim. Presenting reasons could be one way among many of adopting the causal mode in pursuit of getting A to act in a certain way. But the deliberative mode is distinct. When B adopts the deliberative mode, B's reason for discussing A's reasons with him is fundamentally deliberative in nature. She takes talking with A about the reasons he has to be valuable for its own sake because it is a way of recognizing and appreciating his agency, rather than as one way she might influence what he will do (as she might if she were engaging with A in the causal mode).

In §2.2, I argued that sometimes practical co-deliberation is the appropriate way for A and B to interact. It will follow that sometimes B may sincerely try to rationally persuade A not to φ and A may φ anyway. In fact, it may even be that A wouldn't have φ 'd but does so as a direct result of talking with B, despite B's best efforts. Perhaps A is obstinate, or dislikes B, and B's discussion with him results in his reluctance to act the way B recommends.¹² In such a case, because practical co-deliberation was called-for, we should understand B's contribution to A's eventual action in a way that is fundamentally non-causal and non-predictive.

For example, suppose Bonnie engages in practical co-deliberation with Anthony by making a sincere effort to persuade him not to lie by talking with him about why lying is wrong. And suppose that practical co-deliberation is appropriate given their context. Despite their discussion, Anthony goes on to lie anyway. As it happens, Anthony would not have lied if Bonnie had not talked with Anthony about why lying is wrong: Anthony decided to lie to spite Bonnie. So Bonnie causally contributed to Anthony's lying. Even if Bonnie could have predicted that Anthony would lie following their discussion, Bonnie is not complicit in Anthony's lying despite her causal contribution to it. It will be a mistake to criticize Bonnie for Anthony's lying on the grounds that Bonnie did not

¹² Of course, B won't count as engaging in genuine practical co-deliberation if B is aware of A's obstinance and refuses to take it into account, by e.g. "egging A on." This will simply be a different way of manipulating A, perhaps masked in the guise of practical co-deliberation.

do all she could have to prevent Anthony's lying. The fact that co-deliberation was appropriate implies that Anthony and Bonnie were not in a situation in which it was appropriate for Bonnie to try to prevent Anthony's lying by different means. Rather, it was appropriate for Bonnie to adopt a deliberative stance toward Anthony and try to convince him not to lie for the reasons he had not to lie. When the deliberative mode is appropriate, what matters morally is that Bonnie met her deliberative duty, not whether her speech causally contributed to Anthony's lying. Suppose Bonnie had instead encouraged Anthony to lie. This would constitute a failure in her deliberative duty, and so Bonnie would be deliberately complicit if Anthony went on to lie.

My suggestion is that there are two fundamentally different kinds of moral complicity. One kind is somewhat familiar, and I will call it *causal moral complicity*. It is causal in nature and results from a causal contribution or a failure to intervene or prevent A's wrongful φ -ing in a situation in which causal contribution is what is relevant to complicity or intervention was called-for. The second kind is distinct, and it is the *deliberative moral complicity* that I proposed in Chapter 1. Deliberative moral complicity arises when an agent B fails in a deliberative duty she has with respect to A's wrongful φ -ing and A goes on to φ . This deliberative duty failure might be, for example, failing to discuss with A the reasons that A shouldn't φ . The primary focus of the rest of this chapter will be to explain when and why each kind of complicity is on the table when it is and to defend the default entitlement that B has to adopt the deliberative stance toward A.

2.4 Default Entitlement

The default is that B should adopt a deliberative stance toward A and so fulfill her deliberative duty by discussing with A the reasons that A should not φ . In other words, the default is to try to engage in practical co-deliberation. But it is important to recognize that B's entitlement so to act does not rest on a prediction of the sufficient probability of A's acting well conditional on B practically co-

deliberating with him. Instead, adopting a deliberative stance toward other agents is a normative default position. Though there are circumstances in which the default entitlement is not in place, this does not mean that in the standard case our entitlement to adopt a deliberative stance rests on first confirming that such undermining conditions do not hold.¹³ The entitlement is non-predictive in nature. So B is not first to predict how A will act based on different things that B might say or do and then to say or do what is most likely to result in A's performing the right action or avoiding the wrong one. (To do so would already be to adopt a causal stance toward A.) Rather, B is to say what's so—to meet her deliberative duty by co-deliberating with A, by talking with A about the reasons that A should not φ .

But sometimes the default entitlement to adopt the deliberative stance is not in place. What kinds of circumstances undermine the default and require a shift into the causal mode of engagement? I think it is likely to be dependent on the relationship between the parties, the content of the potential wrongdoing, and the context of their interaction. I will not here aim to detail all of the possible undermining conditions, but it seems clear that, for example, entitlement to remain in the deliberative mode is more easily undermined in the case of a parent talking with a child about what to do than it is in the case of two adults talking about what to do. I may permissibly threaten my child with leaving the park if she will not put on her jacket (having already discussed the reasons for wearing one and failing to convince her). And in fact, given our relationship and my responsibilities to my child, I may be required to shift into the causal mode. But I may not make the same threat concerning a jacket to another adult. If I learn that the other adult plans to do something that will result in a great deal of harm to others, though, the default entitlement may be

¹³ This entitlement to adopt a deliberative stance toward A bears structural similarity to our “epistemic entitlement” to rely on our perceptual beliefs (Burge 2003; 2020). Just as in the perceptual case, we can find ourselves in circumstances in which the default entitlement is not in place, but this does not mean that in the standard case our entitlement *rests* on first confirming that those undermining conditions do not hold.

undermined. And if enough is at stake, walking away will be insufficient: intervention may be required. Of course, there will be harder and easier cases: the main idea here is that the default position is an entitlement to adopt the deliberative stance toward others, though this still means talking with them about what to do and why, not acting as though what they do has nothing to do with us.

Notice that if B fails in her deliberative duty in a situation in which it was appropriate for her properly to co-deliberate with A, and A goes on to φ , then B is morally complicit in the thing that happens in the world. B is complicit in A's φ -ing by way of failing in her deliberative duty. She is not complicit in A's φ -ing because of the causal difference she makes or fails to make to A's φ -ing. Rather, she is complicit exactly *because* she fails to meet her deliberative duty. Failing to meet her deliberative duty is constitutive of an instance of moral complicity in A's φ -ing, if A φ 's. On the other hand, since B is entitled to adopt a deliberative stance as a default, the view also implies that if the default is in place, then B will not be complicit in A's φ -ing so long as B meets her deliberative duty, even if B *does* play a causal role in A's φ -ing.¹⁵ So, as we saw above, Bonnie is complicit in Anthony's lying due to her failure to fulfill her deliberative obligation. But if she fulfills that obligation and yet their conversation about lying causally contributes to Anthony's lying, then she will *not* be complicit, despite her causal participation, so long as the default entitlement to act adopt a deliberative stance remains in place.

What if her causal participation in his lying goes beyond simply participating in a conversation about whether to lie? Is her causal complicity really excused just because the default entitlement to adopt a deliberative stance is in place, no matter what she does? Suppose Bonnie

¹⁵ As we will see in Chapter 3, it will sometimes be a complicated matter whether B meets her deliberative duty (and *some* kinds of causal roles may count as such failure).

discusses with Anthony the reasons he shouldn't lie but then participates in his lie, perhaps embellishing his story as he tells it or staying silent as he lies despite knowing he isn't speaking sincerely. Causal participation that extends beyond practical co-deliberation is not part of what it means to engage with another agent in the deliberative mode, and so the default entitlement does not justify this kind of participation in wrongdoing that goes beyond practical co-deliberation. So causal complicity is quite possibly present in these cases. And as we will see in Chapter 3, they may also involve deliberative complicity as well, since contributing to a lying tale may constitute deliberative duty failure.

What is the relationship between the causal mode and causal complicity? If the default entitlement has been undermined and interaction with another agent requires our shifting into the causal mode, we may be required to tailor our conversation with that agent based on what we think is most likely to causally contribute to their avoiding acting wrongly. For example, I may be purposefully cryptic or simply refuse to provide information that would help someone perform a wrongful action. If the default entitlement to adopt a deliberative stance toward the other agent has been undermined, I am at risk for causal complicity in virtue of our conversation in ways that I am not when the entitlement is in place.

It is also worth recognizing *why* there might be undermining conditions that divide our engagement with others into these two sorts of modes. Recall that it is constitutive of adopting a deliberative stance toward another person that you talk with him about what to do and why to do what's to be done. This entails that sometimes such a stance will be made impossible by the stance of the other person. If he is determined to act wrongly and that is apparent, he has closed off the possibility in this case of talking about the reasons together in the way of the valuable interaction that I described above in §2.2. Perhaps he is determined to harm you, and talking about the reasons isn't possible as

you defend yourself. Thinking of the value at stake helps make it more plausible that there will be a point at which B needs to switch to the other mode, to the causal mode. It will not be a calculation, for example, in which B will have to weigh the values at stake and decide that in this case the harms to be done by A outweigh the value of A's free choice. Rather, the valuable sort of interaction simply won't be possible in some cases. A will have made such an interaction impossible.

One advantage of recognizing these two distinct forms of engagement with other agents is that it seems to capture some of the moral complexity we experience when interacting with those committed to wrongdoing. For example, it seems to help capture the internal conflict that B sometimes faces as she tries to decide how to proceed, a conflict that doesn't seem fully explicable unless we recognize that there are two modes of treatment. As such, one small piece of evidence for the idea that there are these two modes of treatment—the deliberative mode and the causal mode—is how it feels from the inside, as the person who is forced to shift from the deliberative mode to the causal mode, when predicting and responding is appropriate. There is a reluctance to treat the other person in this way, even when it is called-for. It is regrettable that we must so treat them, and it is a loss. It isn't the best way to be interacting with that other person, even though it is they who are so determined to do something that is wrong enough to shift us out of the default entitlement to adopt a deliberative stance toward them. And it is striking that in the moment we may even resent the other person for it—resent their forcing us into a position in which we must merely predict what they will do and respond appropriately. As they drag us out of deliberative space and into causal space, we wish they hadn't, even though it is they who we must now treat as something to be managed.¹⁶

¹⁶ This is an interpersonal version of a kind of conflict that has been more widely noticed in the intrapersonal case, especially in the example of Professor Procrastinate (Jackson and Pargetter 1986; Jackson 2014). Professor Procrastinate has been asked to write a book review but knows of himself that if he accepts he will not write it before the deadline, causing problems for the author. By assumption, it would be best if he accepted and wrote, worst if he accepted and missed the deadline, and better than the worst case if he declined, since then someone else would write by the deadline.

2.5 Conclusion

When the default entitlement to adopt a deliberative stance toward A is in place, B understands herself as collaborating with A about what to do. And in so doing, she understands herself neither as causally contributing to A's φ -ing nor as completely irrelevant to and disconnected from A's φ -ing. She sees herself both as relevant to and connected to what A does, not as a matter of causal connection but rather rational connection. B's speech and presentation of reasons are some of the inputs to A's rational deliberation and as such, she is connected to A's eventual action. But the relevant connection is non-causal.

I have suggested that we might think that moral complicity is structured in this way in part because it makes possible a sort of valuable trusting interaction between two agents. More work remains to be done on the content of the co-deliberative duties, which I began to explore in §2.2. I suggested that the co-deliberative duties will be greatly affected by the context of the interaction between the two people and vary based on circumstance, the history of interaction between the two parties, the nature of the values at stake, and what they know about each other. All of these and more, it seems to me, affect the standing permission to talk about the possibilities and reasons with one another, and I will say more about them in later chapters.

Understanding moral complicity in the way that I propose here is promising in its potential explanatory power because it may be able to shed some light on other problems moral complicity raises more generally. For example, the politics of protest can seem mysterious on a causal view. There is a compelling idea that faced with an unjust war, for example, it is my duty to speak out, and

It is controversial what Professor Procrastinate "should do" under these circumstances, given that it can seem inconsistent with his understanding of himself as free that he knows in advance he will not write by the deadline. Declining the review *for the reason that he knows he won't write it*, despite the fact that he *could* write it and only *won't* write it, can seem like an odd form of "self-management". Thanks to Pamela Hieronymi for introducing me to this example and for helping discussions surrounding it. For interpersonal analyses that share my worries about this kind of "management" of others, see Manne (2014) and Basu (2019), especially Basu's example of what is wrong with the way Sherlock Holmes treats others.

it can seem those who remain silent are morally complicit. But this idea can be hard to make sense of in causal terms. My decision to speak out against the war will make no difference to whether or not the war continues. And yet the intuition that silence can be complicity stands.

This view of complicity holds the promise of recovering the judgment that to remain silent is to be complicit, even when one's speaking up will make no difference. To speak out is to argue, by appeal to moral reasons, that the people making the war should stop. To speak out is a way of meeting one's deliberative duty. And when I fail to speak out, I do not fail to do one more thing I might have done that might have made a causal difference to whether or not the war continued. Rather, I fail to do my deliberative duty. And so, I can count as sharing responsibility for the wrongful thing despite the fact that my own stance makes no appreciable causal impact.

To push the example just a bit further, the view might also help to explain the intuition that my duty to speak up is strongest in democratic or putatively democratic conditions. These are not conditions where my protest stance is likely to make a difference, but they are conditions under which people in power have made a show of asking my opinion, thus possibly triggering a higher grade of deliberative duty than applies under outright tyranny.

As the protest example helps to bring out, in further work I would like to try to use the structure of complicity I defend here to try to make sense of the more general landscape of complicity, including outside of the very narrow set of cases in which it's relatively clear that the deliberative duties have been triggered because someone has asked me what to do.

Chapter 3: Silence as Complicity and Action as Silence

3.1 Introduction

Flagrant injustices cry out for action, such as protest, whistleblowing, even civil disobedience. Some protestors might march under signs that say, “Silence is Complicity.” A whistleblower might attest that he couldn’t remain silent any longer; he felt he had to speak up or else become complicit. And these ideas resonate with us: it seems that at least sometimes remaining silent in the face of injustice amounts to complicity in the injustice. What philosophy owes is an account of how this could be so.

If we take the slogan “Silence is Complicity” seriously, it seems to claim that the silence itself *constitutes* the complicity. So we are not just interested in an account of moral complicity that is extensionally adequate—that captures cases of silence as among the cases of complicity (though of course extensional adequacy is important). We are interested in something more explanatorily powerful, an account of moral complicity that can capture how it could be that the silence *itself* amounts to complicity. In other words, explaining how it could be that silence *constitutes* complicity is a desideratum on the best philosophical account of moral complicity.

What might such an account look like? The deliberative view of moral complicity holds promise: silence itself might constitute a failure in one’s duties concerning others’ practical deliberations and so constitute complicity. We will certainly want to know more, such as when and under what conditions silence counts as such failure. The first main goal of this chapter is to explain how the deliberative view of moral complicity can vindicate the intuition that silence sometimes constitutes complicity.

The deliberative view of moral complicity faces an important and distinctive problem, though. It does not seem correct to say that our *speech* (and so also our non-speech through our silence) is all that can make us complicit. Often, our complicity seems to have nothing whatsoever

to do with deliberation. Rather, it seems that in many cases our *action* (or even lack of action) is what makes us complicit. Consider an example. Suppose Betsy drives Adam to a bank so that Adam can rob it. Betsy intuitively seems complicit in Adam's bank-robbing, but can a deliberative view explain why this is so? Even if Betsy had told Adam that robbing a bank is a wrongful action, it still seems like her *driving* Adam to the bank makes her complicit in Adam's bank-robbing. The deliberative account owes an explanation of how the *driving* can constitute complicity, if it is to apply to more than just a very narrow set of cases of complicity. The second main goal of this chapter is to show that the deliberative view of moral complicity can apply to and explain cases it may (at first pass) seem to miss, such as cases of action.

If the argument of this chapter is successful, it provides a distinctive reason to prefer the deliberative account of moral complicity: it can capture the silence itself as complicity *because* it is silence. And further, by turning our attention toward our deliberative obligations, a deliberative account can help us to see *when* and *why* silence amounts to complicity.

3.2 "Silence is Complicity!": A slogan that comes with a puzzle

Silence in the face of wrongdoing sometimes constitutes complicity in that wrongdoing. Here my aim is not to argue on behalf of this idea but instead to explain why it raises a puzzle for philosophical accounts of moral complicity. Consider a couple of examples:

Silent Friend

Bernadette and Alice are friends. Alice tells Bernadette that she is cheating on her spouse and lying to them about it. Bernadette says nothing to Alice about why Alice's actions are wrong.

Silent Board Member

Benedict sits on the board of a major corporation. The corporation is engaged in activities that are seriously wrong, and Benedict is aware of those activities. He does not say anything in opposition to those wrongful activities to either his fellow board members, regulatory agencies, nor the press.

In these cases, Bernadette and Benedict seem complicit in wrongful actions in virtue of their *silence*.

But this raises a philosophical puzzle: how could they be complicit when they're not doing anything?

And their not doing anything seems precisely to be the problem. But how could the not doing anything *itself* constitute the complicity?¹ It will count in favor of an account of moral complicity if that account can provide a solution to this puzzle and explain in a satisfying way how silence can constitute moral complicity. How might such a story go?

3.3 Silence as Problematic Deliberative Contribution

A turn to the deliberative view of moral complicity seems promising. Recall that according to such a view, an agent is complicit in wrongdoing not in virtue of her causal contribution—recall the default entitlement to act trustingly toward other agents—nor her intentional participation in that wrongdoing but rather because of a failure to fulfill a deliberative obligation she had with respect to that wrongdoing.

Deliberative View of Moral Complicity:

B's failure in her deliberative duty regarding A's φ -ing is *sufficient for* and *explains* B's moral complicity in A's wrongful φ -ing, if A φ 's.

The deliberative view holds that at least some cases of complicity can be explained in virtue of deliberative duty failure.² It contends that in some cases of complicity, an agent B is complicit in some wrongful action of another agent A's *when* and *because* B fails in B's deliberative obligation to talk with A about what's to be done and why. For example, B might fail in her deliberative obligation with respect to A's wrongful φ -ing by encouraging A to φ . But notice that B might *also* fail by staying silent in the face of A's wrongdoing when the context of her situation makes speaking

¹ We might notice that other views of moral complicity, such as views based in causal contribution or intentional participation, seem to encounter difficulty explaining how doing nothing could constitute complicity. If omissions count as causes, then a view based in causal contribution may be able to explain silences as complicity, but it will be overinclusive unless a further factor distinguishes those omissions that count as relevant and those that do not. As for intentional participation, it is hard to see how cases of silence could be understood as instances of participating intentionally when precisely what is happening is a failure to get involved.

² Recall the default entitlement structure of Chapter 2, which leaves open the possibility of causal complicity in other cases.

up called-for. This means that the deliberative view of complicity has helpful resources to offer in explaining the slogan “Silence is Complicity” and the moral force it carries. Bernadette is complicit in Alice’s wrongful treatment of her spouse because once Alice makes the moral content of her behavior salient between them Bernadette has an obligation to speak out. The silence itself *is* Bernadette’s complicity: she has an obligation to speak up, and since she does not, she is complicit in Alice’s wrongdoing. Benedict is complicit in the corporation’s wrongful activities because his role as a board member comes with an obligation to say something about why the corporation should not be involved in those activities. Again, the silence itself *is* Benedict’s complicity: he has an obligation to say something, and since he does not, he is complicit in the wrongful activities.

Thus, the deliberative view of moral complicity holds the promise of explaining our philosophical puzzle concerning the complicity of silence in a satisfactory way. How could doing nothing make one complicit? It can make one complicit because doing nothing, when one has a duty to speak up, violates that duty, constituting problematic deliberative contribution and so moral complicity.

This raises an immediate question: when is speaking up required? It may be that I benefit from an unjust scheme, and my benefitting explains why I am required to speak though others are not so required. Or perhaps someone has asked me for advice about what to do. I may also be required to speak up if the context makes a particular content salient between me and another agent. Of course, we will disagree about the precise contours of the deliberative duties, but substantive disagreement about what content is salient between which agents is independent of the structural contention I am defending here that deliberative duty failure opens one up to moral complicity.³

³ We may, for example, disagree about individual cases and whether a particular content is salient in a particular context and between two particular agents. Scanlon makes a similar point about objections to his arguments defending contractualism. An objector might disagree with Scanlon’s evaluation of what principles people could reasonably reject, but this would not be an objection to contractualism as such but only to Scanlon’s evaluation of a particular case. Scanlon calls this kind of disagreement a “substantive disagreement,” and I follow his use of the phrase (2000, 186, 205, 238).

Sometimes agents fail in their deliberative duties by failing to speak up, and the deliberative view of moral complicity provides helpful resources to determine whether a deliberative duty is in place between two agents with respect to a particular content, meaning that it provides some helpful resources to make progress on these disagreements.⁴

Sometimes, our duty to speak up is grounded in another of our deliberative duties—our duty of care with respect to the beliefs of others.⁵ This duty also sometimes requires making corrections to others’ speech.⁶ As a whole, the duty requires that we take reasonable care to ensure that others’ false beliefs are not formed nor confirmed because of our conduct, including our behavior, speech, and omission.⁷ What does reasonable care consist in? I have in mind something similar to a reasonable person standard as it is understood in legal theory.⁸

A few examples might help. As a start, it is easy to see that some cases of intentional deception violate this duty of care.⁹ If Bob spills coffee on Sue’s laptop and wants to avoid paying to fix or replace it, he might dip their cat’s paws in cold coffee with the aim of misleading Sue about the culprit.¹⁰ This is a clear violation of Bob’s duty of due care with respect to Sue’s mental contents.

But one could also violate this duty unintentionally and unknowingly by, e.g., being unreasonably uncautious in one’s speech or action, as when one says something *false* that one

⁴ I will also say more about these details in Chapter 4.

⁵ As I understand this duty, it does not rule out the possibility of permissible deception in cases of e.g. non-lying deception intentionally deployed to protect one’s reasonable sphere of privacy. See Shiffrin (2019, 79).

⁶ For example, in some contexts agents have a “duty to object” when someone else says something untrue. Lackey argues that such a duty has both moral and epistemic components (2020). I take no stand on the debate concerning the distinction between moral duties governing the epistemic domain and “pure” epistemic duties here. For an argument that the epistemic duties are intrapersonal and strictly doxastic, see Wrenn (2007).

⁷ For a similar view in the context of the wrong of deception, see Shiffrin (2019).

⁸ See (among others) Holmes (1909, 54). Of course, care and caution in our application of such a standard would be crucial. For examples of important criticisms of the reasonable person standard in practice, see (among others) Peterson (1999), Susskind (1993), and Carpiello (2001).

⁹ For a discussion of intentional deception and unintentional deception as well as an argument that some cases of intentional deception are permissible, see Shiffrin (2019).

¹⁰ This example is due to O’Neil, who uses it as an exemplar (though an admittedly contrived one) of what he calls “covert deception” (2012, 302).

believes or when one unintentionally deceives another, if the deception is unreasonable.¹¹ Or one might fail in the duty by unconsciously allowing someone to draw a false inference based on one's speech (though of course one is not responsible for *any* false inference that a listener might draw on the basis of one's speech). This duty of care is both content and context dependent: what level of care counts as reasonable varies based on the context of my speech and based on what I say. The general idea is that one is not "off-the-hook" with respect to the beliefs of others: minimally, we must exercise reasonable care to ensure we aren't the source of their epistemic mistakes.

Finally, it is worth noting that one can violate this duty with respect to another person even if that person does not in fact form nor have confirmed a false belief on the basis of one's conduct.¹² This is because one might "get lucky".¹³ The other person might do further research, or just not take one seriously. And yet in such a case the duty violation stands.¹⁴

3.4 But what about the *driving*?

Deliberative complicity might seem like a fine way to capture *some* cases of complicity, but its scope can seem quite narrow. After all, in many cases, what matters seems to be what the complicit agent *did*, not some discussion she had or failed to have with the agent who performed the wrongful conduct. And if the scope of deliberative complicity is limited enough, it may even fail as an explanation of *silence* as complicity since it can seem like there isn't really a unified phenomenon of interest here: complicity through deliberative duty failure might just be a completely different concept than the moral complicity we usually have in mind.

¹¹ Here by "unknowingly" I mean to include even cases in which the failure is not consciously accessible to the person who has failed in her duty. She may not know the truth or may not know that the other person is likely to form or have confirmed a false belief on the basis of her conduct, but neither of these is dispositive with respect to her duty failure.

¹² Notice that this isn't true of deception since "deceive" is a success word. It implies that the person who has been deceived (deceivee?) has *in fact* been deceived, so they must have formed a false belief or had one confirmed.

¹³ I have in mind the same sense of "lucky" as Nagel means to ascribe to the negligent driver who despite his negligence does not hurt anyone (1979).

¹⁴ For an argument that one fails to fulfill one's duty even if one gets lucky in this sort of way, see Herman (2019, 11).

Let's start with an example. Consider a case often thought to be a core case of complicity: Betsy is the getaway driver at a bank robbery; Adam robs the bank. Each is aware of what the other is doing and why. In this case, Adam is the primary agent, and Betsy is the secondary agent.¹⁵ Betsy seems complicit in Adam's bank-robbing. Why? One intuitive answer is that Betsy causally contributes to Adam's bank-robbing. But, as I argued in Chapter 1, causal contribution is unsatisfying as a necessary condition on complicity and cannot possibly be a sufficient condition on complicity. Nonetheless, a causal theory of complicity seems well-positioned to capture Betsy as complicit in Adam's bank-robbing and to capture similar cases because it provides a simple, plausible, and compelling explanation of Betsy's complicity: Betsy is complicit in Adam's bank-robbing because her *driving Adam* constitutes a causal contribution to his bank-robbing.

But a deliberative view seems less well-positioned to explain this sort of case. Betsy seems complicit not because of some contribution she had with Adam about whether robbing a bank is morally permissible but rather because of her *act* of driving. But her act of driving isn't an obvious failure in any deliberative obligation she seems to have. We might even imagine that Betsy in fact *had* an extended conversation with Adam in which she tried to convince him not to rob the bank. And yet if she drives the car, her driving seems to ground her complicity, regardless of whether their discussion fulfills her deliberative duty. The deliberative view of complicity does not seem well-positioned to explain this sort of moral complicity, the kind that is grounded in an agent's action rather than her speech.

I think that the view can explain such cases, contrary to appearances. It can do so by pointing to the ways in which our *actions themselves* communicate our stances on moral issues.

¹⁵ I follow legal discussions of criminal complicity and refer to complicit agents as *secondary* wrongdoers, in contrast with the primary wrongdoers who commit the primary wrong (the wrong in which the secondary wrongdoer is complicit). Discussions of criminal complicity also sometimes call the primary wrongdoer the *principal* and the secondary wrongdoer the *accomplice*. See e.g. Gardner (2007), among others.

3.5 Acting-as-Though: Action as Silence

It is part of our common practice to impute stances to agents based not only on their speech but also based on their actions. In the course of trying to decide who to vote for, we might try to ascertain the candidates' positions. Such an effort involves more than looking only at what they say: their actions are also relevant. If one candidate *says* she cares about income inequality but has regularly voted against worker protections and in favor of tax cuts to the wealthy, we understand her as opposed to governmental intervention to combat income inequality despite her claims to the contrary. Similarly, we might be skeptical that someone is a vegetarian, even if he claims he is, if he regularly eats meat. I might doubt a friend who attests she values my friendship if she never calls nor writes and regularly declines invitations to spend time together. In each case, what the agent says isn't the final word on what we understand their take on the world to be.

The point here is a rather small one: we impute stances to agents on the basis of more than just their direct speech. We impute stances to them based on their non-speech conduct, including their actions, omissions, and the implications of their direct speech. I think we should understand agents' conduct as expressive, as expressing their take on the world. We might say that when agents act they *act-as-though* certain things are true, meaning that they *express* that those things are true.¹⁶ For example, an agent may *act-as-though* φ -ing is permissible when she φ 's.¹⁷ It is important to note that this does *not* mean that she *believes* that φ -ing is permissible. She may, of course, but she also may not. We can (and do!) act in ways that we believe are wrong. We treat others poorly, miss deadlines,

¹⁶ But they needn't *believe* that those things are true. Just as agents can say what they do not believe, they can also express through action what they do not believe. For similar views about the objective meaning of actions (and agreement that such meaning does not depend on the actor's beliefs) see Helmreich (2015, 88–90) and Anderson and Pildes (2000, 1508). I use “stance” differently from the way Helmreich does, but his “behaving toward S as though P” is quite close to my “*acting-as-though* P”. (For Helmreich, it is essential to a stance that one's action is undertaken *for the reason that* P, where P is the normative claim that one accepts (2015, 97-98). My use of “stance” does not require that the agent act for any particular reason at all.)

¹⁷ Callard makes a similar point: “But I, who did it, evidently saw it as a perfectly fine thing to do, having judged the action to be a good thing for me to do” (2020).

fail to recognize needs that we ought to notice and address. And sometimes we do these things even though we would agree that we shouldn't do them. But our actions *themselves*, independently or our speech and independently of our beliefs, express our stance on issues: when we do what we shouldn't, we express that what we shouldn't do is permissible, for us, under those circumstances.

When a person says that P, she expresses her stance that P is true, regardless of whether she *in fact* believes that P.¹⁸ Of course, agents can (and do!) say what they do not believe. But if the context is the usual one then saying that P amounts to an expression that P is true.¹⁹ For example, suppose I say that the sky is blue today but in fact believe it is gray. This is perfectly possible: we can say what we don't take to be the case. Nevertheless, I have *expressed* the stance that the sky is blue. If you reply that you think my assessment mistaken, that today it is cloudy and thus the sky is gray, I cannot respond to you by saying, "Oh, yes, that's what I said!" Once I've said that the sky is blue, you too, can attest to what it is that I have said.

3.6 Acting-as-Though and Duty Violation

How does this connect with the duty of due care to the contents of others' minds? If an agent *acts-as-though* φ -ing is permissible even though it is not, such expression can constitute a failure to exercise reasonable care that others' false beliefs are not formed nor confirmed because of her conduct.

¹⁸ I take this as a starting point and do not argue for it here. For an argument in this claim's defense and a discussion of how this might work, see Moran (2018, especially Chapter 3). He says that a speaker "expresses her belief that P", but I avoid the use of "belief" since I think the one can express what one does not believe and so "belief" can be misleading in this context. (Moran also does not think that the expression is constrained by the actual beliefs of the speaker.)

¹⁹ In some contexts, saying P does not amount to a commitment to P's truth. One such context is a "justified suspended context," which the normative presumption that a speaker only says what she *believes* to be true is suspended, such as it is in the context of play-acting, improv performances, or telling a joke. It is crucial to such contexts that the fact and justification of the suspension are publicly accessible (Shiffrin 2014, 16). Another such context is one in which the normative presumption that a speaker only say what *is in fact* true is suspended, such as when she says, "don't count on me, but P" or, "I'm not sure, but P."

Employing this concept of *acting-as-though*, the deliberative view can explain and ground the complicity of agents not just when they speak in favor of a wrongdoer's action but also when they support that action with their own action or omission. The deliberative view can employ the resources of the *acting-as-though* theory in order to get the intuitively correct result about the *fact* of complicity in the relevant cases. And, in such cases, the explanation of complicity will remain fundamentally deliberative. The reason Betsy is complicit in Adam's bank-robbing is that her driving the getaway care deliberately supports Adam's bank-robbing. And it is the deliberative support the driving provides, not its causal contribution, that explains her complicity. Her driving the car expresses her stance not only that her driving is permissible but also that Adam's bank-robbing is permissible.

This raises the question: how, when, and why does B's action express that A's action is permissible? How do we move from her expression of a stance on driving to an expression of a stance on robbing? This is a bit complicated, so we will consider it in detail in the next section.

3.7 A Worry for *Acting-as-though*: Is there circularity here?

As we have seen, since Betsy's driving causally contributes to Adam's bank-robbing, a view based in causal contribution holds her morally complicit *in virtue of her driving*, and this seems like the right result. But what will a deliberative view have to say? Suppose Betsy has never argued in favor of bank-robbing.²⁰ A deliberative view may seem more appropriate in holding Congressmen responsible for the effects of laws the passage of which they defend with grand rhetoric, since their deliberative support is clear in such cases. But if Betsy makes no such argument, what is her problematic deliberative contribution supposed to consist in?

²⁰ Note that even if she *had* argued in favor of bank-robbing the deliberative view still seems to face a difficulty: she will be complicit in virtue of her so arguing but *not* in virtue of her driving.

As a first step, we might say that in virtue of her driving, Betsy *acts-as-though* Adam's bank-robbing is permissible. If this is right, then she violates the duty of due care she has to express the truth—that bank-robbing is wrong. And then it is in virtue of Betsy's driving, because her driving counts as *acting-as-though* and so as a failure in her duty of due care, that the deliberative view holds her complicit.

But is this first step right? *Why* does Betsy's driving constitute *acting-as-though* Adam's bank-robbing is permissible? And can the deliberative view adequately address this question without a worrisome sort of circularity?

The question we are faced with is whether what Betsy counts as expressing through her action (what we might call the *content* of her *acting-as-though*) can be determined and grounded without *already* assuming she is complicit in Adam's bank-robbing. For suppose we thought that the reason Betsy's driving amounts to *acting-as-though* Adam's bank-robbing is permissible were because driving Adam makes Betsy complicit in the bank-robbing and so for *that reason* amounts to an *acting-as-though* it is permissible to rob a bank. Then the deliberative view would provide a circular explanation of Betsy's complicity, since in order to explain her complicity we appealed to her *acting-as-though* but in order to explain her *acting-as-though* we appealed to her complicity.

What the deliberative view owes, then, if it is to cover such cases on independent grounds, is an explanation of why Betsy's driving constitutes *acting-as-though* Adam's bank-robbing is permissible that isn't *itself* grounded in or based on an assumption of Betsy's complicity.²¹

Here is a try: Betsy had to settle the question of whether to drive. This means that she ought to have considered whether driving was permissible. And this means that it made sense for Betsy to

²¹ By independent grounds, I mean grounds that would allow the deliberative view to explain Betsy's complicity in deliberative terms without grounding the explanation of deliberative complicity on complicity of another sort. As an alternative, the deliberative view might ground Betsy's deliberative complicity on her causal complicity, if a causal account could be made to work as an explanation of at least some kinds of complicity, including the Betsy's case.

ask whether what she was driving Adam to do was permissible for him to do. Given her involvement in his activity—her supporting it by her driving—the permissibility of *his* activity is a consideration that bears on the question *for her* of whether to drive.²² When Betsy drives, she *acts-as-though* she has settled the question of whether to drive is permissible in the affirmative. And because the permissibility of Adam’s activity bears on whether her driving is permissible, she also *acts-as-though* his activity is permissible.

Why does it make sense for Betsy to inquire about the permissibility of Adam’s action? To see this, it helps to think about a different kind of case in which what Adam is being driven to do is permissible or even good, perhaps meet his daughter at the airport. In this kind of case, notice that the fact that Betsy’s driving is a *helping* Adam to meet his daughter is a reason to drive. That is, it counts in favor of the driving that the driving is a helping. But this is only the case because what Adam is up to is permissible. If we return to the robbing case, it *could* be that the fact that Betsy’s driving is a helping to rob is a reason to drive. But this could only be true if it were permissible to rob, since it won’t count as a reason to drive that the driving is a helping to do something that is impermissible. So driving that is reasonably taken as an *acting-as-though* helping were a reason to drive is thus *acting-as-though* robbing were permissible. And *that* is why the driving is an *acting-as-though* robbing is permissible.

In other words, it is not the *very fact* that the bank-robbing is impermissible that makes the driving impermissible. The moral landscape is much more complicated. The driving constitutes complicity in the bank-robbing because the driving is a violation of the deliberative duty to take reasonable care that others’ false beliefs are not formed nor confirmed on the basis of one’s conduct. The driving is a failure with respect to that duty because it amounts to *acting-as-though* robbing is permissible. And it amounts to *acting-as-though* robbing is permissible because of the

²² Because she is lending her agency to his activity, the permissibility of his activity bears on her deliberative question.

deliberative relevance of the permissibility of robbing to the decision about whether to drive. And finally, the permissibility of robbing is deliberatively relevant to the decision about whether to drive because driving could reasonably be understood as helping, were helping something one had reason to do in this case.

This is a significant step toward understanding complicity, but it also brings with it important questions. *When* is it that the permissibility of *another* agent's action bears on my practical question of what to do? A full exploration of this question would take us outside the scope of this dissertation, and I hope to think about it more in future work, as it connects with our interest here in interpersonal deliberative responsibility. Some core cases seem obvious; some cases at the margins seem more difficult. If Kyle has asked me to endorse or praise his action, its permissibility is relevant my practical question of whether to praise or endorse. On the other hand, the permissibility of Yvette's expensive car purchase is not relevant to my practical question of which charity I will support this month. But it may be hard to say whether the permissibility of Ted's plan to spend a lot on his dinner bears on my question of whether to support his night out by watching his children for the evening.

Cases in which the message expressed by the potentially complicit agent's action is ambiguous raise their own sorts of difficulties. For example, suppose meat-eating is wrong, Betsy believes it is wrong, and she drives Adam to pick up his lunch that includes meat. Is she complicit in his meat-eating? According to the deliberative view, this comes down to whether or not her driving him to pick up his lunch constitutes a failure in her deliberative duty. And it seems like her driving may or may not constitute such a failure, depending on the case and some of its features, such as the background of their deliberative relationship. Are they close friends? Are they in the habit of discussing moral issues with one another? If they *do* have a standard of such discussions, then an omission here on Betsy's part may constitute a failure in her deliberative duty. But if they

are not close, and they don't normally have such discussions, I expect that just a simple driving wouldn't constitute *acting-as-though* and so wouldn't constitute failure in a deliberative duty. Why not, given that Betsy's driving Adam in the bank-robbing case *does* constitute *acting-as-though* and so deliberative duty failure? I think it makes a difference that the message communicated by Betsy's driving is more ambiguous in the lunch case than in the bank-robbing case. It is difficult to see what her driving in the bank-robbing case could be other than a helping to rob.²³ But in the lunch case, her driving needn't be a helping to eat meat; it might just be a helping to eat lunch. And perhaps that ambiguity is sufficient to avoid deliberative duty failure.

More difficult still are cases in which the stance expressed by the potentially complicit agent's action is in tension with the stance expressed in her direct speech. For example, we might imagine that Betsy drives Adam to the bank-robbing but does so while talking to him the whole drive about all the reasons bank-robbing is wrong. It seems to me that Betsy is still complicit: the stance expressed by her action seems to belie her speech. But it isn't clear that this will always be true. For example, if Betsy drives Adam to pick up his lunch but does so while talking to them the whole drive about all the reasons eating meat is wrong, she does not seem to me to be complicit in his wrongful meat-eating. If I am right, what separates the cases? When and why does the stance expressed by the action dominate that expressed in speech? I think it is hard to say but that usually the stance expressed by the action wins the day. When it does not, I suspect that it is because the stance expressed by the action is itself ambiguous and so the speech is able to play a clarifying role. For example, Betsy's discussion of the wrongs of meat-eating clarifies that her driving is a helping Adam eat lunch and not a helping Adam eat meat. But these mixed cases are hard, and I don't think

²³ Though perhaps even in the bank-robbing case things are different if there *is* an obvious competing explanation that makes the meaning more ambiguous. For example, perhaps the driving does *not* count as *acting-as-though* robbing is permissible if Betsy is explicit that she is driving for another reason, and that reason is significant enough: "Adam, I think that what you are doing is wrong. But you are a dangerous driver, and I am worried that an innocent bystander will get hurt if you drive. So I am going to drive."

I've settled here how all of them turn out. In future work I hope to think more explicitly about whether the stance expressed by an agent's speech might ever trump the stance expressed by her action.

Allow me to address a lingering worry one might harbor about the driving case. Even if my argument that the permissibility of the driving is deliberatively relevant to the question of whether to drive is convincing, it might still seem rather significant and important that not only did Betsy *act-as-though* robbing were permissible, she also *participated in* a robbery. In fact, it might even seem as though if the deliberative view is right and she has failed in her deliberative duty, that is a fine but significantly incomplete explanation. Then Betsy has done *two* things wrong: (1) fail in her deliberative duty and (2) participate in a wrongful action. And isn't participating in a wrongful action worth worrying about, and perhaps even more important to complicity than deliberative duty failure?

Perhaps: the argument here leaves that possibility open. I have argued that deliberative duty failure is sufficient for complicity, so it could be that participation (perhaps intentional participation) is another way that one can be complicit and even that there are further ways—perhaps some sort of causal relation that results in complicity (when the default entitlement to act trustingly discussed in Chapter 2 is no longer in place). Or it could be that the causal relation or the intentional participation themselves ground *deliberative* complicity because they represent another kind of connection between the driving and the robbing that could independently ground the content of the driver's *acting-as-though*. (So, for example, Betsy counts as *acting-as-though* bank-robbing is permissible because she intentionally participates in bank-robbing.)

But I will offer one last reason for thinking that the deliberative view is distinctively promising. Suppose we are worried about Betsy's participation but a bit hesitant to call her participation complicity. Perhaps instead it strikes us as *cooperation*, where cooperation is a moral

concept distinct from complicity. The deliberative view allows us to draw this kind of distinction, as between *deliberative complicity* and *cooperation*. This gives it an advantage over views based on intentional participation and causal contribution, both of which are unable to draw such a distinction.²⁶

3.8 Protest

The deliberative view's explanation of silence as complicity faces a remaining worry. Most of the examples I have considered involve one-on-one conversations or actions between two people. In such cases, it is clear that the primary agent hears what his conversational partner says or notices what she does. But a protest may not be noticed or watched by the person or people engaged in the wrongful activity. If the message is not received, why should we agree that the protestor has fulfilled a deliberative duty? And why should we think that the person who stays home *fails* in her deliberative duty, if her protesting would not even have been noticed by the agents who matter? Protest seems to be a fundamentally different sort of communicative activity than one-on-one conversation, and so it can be difficult to see how it fits cleanly into a deliberative model that has as its core cases small interactions between two agents.

Protest as communication and the question of just what it communicates is an interesting question all its own, and I will not attempt to do it full justice here. But let me say a bit about why I think the deliberative view *can* capture cases like protest. First, it is a mark in its favor that the deliberative model does not require causal connection. So the fact that the protest escapes the notice of the primary agent or agents does not undermine the possibility of deliberative complicity. What will matter is whether the protestors, in protesting, have fulfilled their deliberative duty. Why

²⁶ I leave it open whether moral complicity might be a large category with cooperation as a sub-category. Even if that is the case, the two moral concepts are distinct.

should we think, though, that they have fulfilled their deliberative duty to someone if he does not even notice their communication, let alone understand it fully? In protesting, they have clearly expressed their stance that the injustice is wrong. They have taken steps to announce their stance and so provided the opportunity, should the primary agent avail himself of it, to communicate the moral truth to him. Even if he fails to take it up, they have fulfilled their deliberative duty.

3.9 Conclusion

The deliberative view of moral complicity holds the promise of capturing and explaining the pre-theoretic intuition that silence sometimes *constitutes* complicity. This is a mark in its favor. But it faces a difficult problem of its own, one that doesn't affect views based in causation or intentional participation: it does not seem obvious how the view can capture *actions*—and not just speech or speech omission—as cases of complicity. In this chapter, I have argued that thinking about the ways in which we attribute expressions of stances to agents based on their actions can help us to explain how the deliberative view of moral complicity can indeed capture actions (and non-actions) as instances of complicity.

In the next (and final) chapter, I turn to some questions remaining for the deliberative account of moral complicity, including the extent of its scope.

Chapter 4: The Scope and Explanation of Deliberative Complicity

4.1 Introduction

A philosophical account of moral complicity must do (at least) two things: pick out *what* an agent is morally complicit in and *explain why that* is what she is complicit in. For example, views of moral complicity based on intentional participation (such as the one discussed in Chapter 1) seem to have a satisfying answer to both of these questions. *What* is an agent morally complicit in? Those projects in which she has intentionally participated. *Why* are *those* the projects in which she is morally complicit? Because of her intentional participation in them.

By contrast, the deliberative conception of moral complicity does not seem to have a clear and obvious way of drawing a boundary between the actions in which an agent would be complicit and that actions in which she would not. It is tempting to say that what an agent is complicit in is determined by the deliberative duties in which the agent has failed. In other words, the deliberative duties determine the *what* of complicity. In a sense, this simple answer is the right one. But in simply pushing the question back a level, the simple answer leaves many of the details open. If Brianna fails in a deliberative duty today with respect to Agatha's cheating, is she morally complicit in all of the cheating Agatha does for the rest of her life? Is there a way for her to correct her failure and thus avoid further complicity, or is she just stuck forever? If she *can* correct her failure, how would she do so? Does she need to call everyone with respect to whom she has failed in a deliberative duty and let them know that she was mistaken and the actual moral truth is X?

In this chapter, I aim to explain and make progress on this "problem of scope" that faces the deliberative conception of moral complicity. We can understand the problem of scope as asking what set of potential actions an agent would be complicit in, were they to occur. I do not aspire to answer all of the difficult questions nor solve every borderline case, but I do hope to take some first steps toward a solution to the problem on behalf of the deliberative conception.

4.2 Duty Failure and Moral Complicity: Facing Questions of Scope

Recall our Agatha and her roommate Brianna from the Introduction. Agatha tells Brianna that she plans to cheat on her upcoming midterm exam, and Brianna nods along as Agatha shares her plan. Though she hopes that Agatha will not cheat, she does not tell Agatha of this hope nor does she raise for Agatha's consideration any of the reasons that cheating is wrong. Agatha goes on to cheat, and according to the deliberative view of moral complicity, Brianna is morally complicit in Agatha's wrongful cheating. Why? Brianna is complicit because she failed in a deliberative obligation she had with respect to Agatha's cheating. Once Agatha raised between them her plan to cheat, her cheating became deliberatively relevant, and Brianna should have spoken against it.

What the deliberative view of complicity offers here is not a new or surprising result about *whether* Brianna is complicit but rather a way of understanding what has happened that can *explain* her complicity. Brianna is not complicit because she caused Agatha's cheating nor because she intended it. And yet she seems complicit anyway. As we saw in Chapter 1, the deliberative view of complicity can explain why she is complicit in the absence of these other conditions.

We might ask why Brianna's duty failure makes her complicit in Agatha's cheating. Recall from §1.6 the nature of the connection between Brianna's duty failure and her complicity: because Brianna failed in a deliberative duty she had with respect to Agatha's cheating, she is complicit in that cheating. Put another way, it matters that Brianna failed in her deliberative duty with respect to (1) Agatha and (2) cheating. If Brianna had instead failed in her deliberative duty with respect to, say, Justin's promise-breaking, she would not be complicit in Agatha's cheating (assuming she met the deliberative duties she had surrounding Agatha's cheating). So here we have the beginning of an answer to our scope question, a start at drawing a boundary around the *what* of Brianna's complicity. Brianna is complicit in Agatha's cheating because she has failed in her deliberative duty regarding that cheating. But she is not complicit in Justin's promise-breaking because she has not failed in her

deliberative duty with respect to Justin's promise-breaking. Agatha's cheating is inside the boundary: Justin's promise-breaking is out.

4.3 Determining Complicity's Scope: Benefits of Intentional Participation

But this is only a first step. How far does Brianna's complicity extend, according to the deliberative view? Suppose Agatha also cheats on her final exam. Is Brianna morally complicit in this further instance of cheating? If they have not revisited their discussion regarding the permissibility of cheating, then Brianna has not remained silent in a conversation with Agatha in which they discussed cheating on Brianna's *final* exam, at least not *in particular*. Does this matter? Suppose the final exam takes place a few weeks after their conversation about cheating and its moral status. Does the timing matter? Are things different if Agatha cheats on a *midterm* the following year, since their conversation surrounded her midterm exam? What about the year after that? What if Agatha convinces Edward to cheat on his midterm exam, repeating the same reasoning she rehearsed in front of Brianna? Is Brianna then complicit in *Edward's* cheating, since she also plays a role in *his* chain of moral reasoning? And what if Edward goes on to repeat the same reasoning to someone else? Or if Cynthia *overhears* Brianna and Agatha's (original) conversation and goes on to cheat on *his* midterm? Is Brianna morally complicit in *his* cheating even though she wasn't speaking to him?

Notice that these questions are particularly pressing for the deliberative view of complicity, since some alternative conceptions of moral complicity seem well-positioned to address them. In particular, the intentional participation view seems to have the resources to draw clear distinctions along the lines of scope. It entails that the complicit agent is complicit in what she participates in intentionally. So if Brianna helped Agatha come up with her cheating plan, intending that Agatha would cheat, Brianna would be complicit in her cheating. But she won't be complicit in the cheating of the person who overhears (unless, perhaps, she is aware of his listening and carries on the

conversation with the intention that he, too, will cheat). Nor will she be complicit in later instances of Agatha's cheating (unless she also intends those). But the deliberative view cannot rely on an agent's intentions to draw these kinds of distinctions. After all, moral complicity, according to the deliberative view, expressly does not require any intention at all. One could fail in a deliberative duty that one had with respect to an agent and a wrongful action on the part of that agent even if one didn't realize one had such a duty or wasn't aware that the agent intended to perform the relevant wrongdoing.

What, then, does the deliberative view of moral complicity have to offer by way of answering these scope questions?

4.4 Deliberative Duty Failure: Direct Speech

Let us begin with the question of *in whose* actions the complicit agent (B) is complicit. First, we can recognize that the person with whom B engages in a conversation is particularly significant. We saw in Chapter 2 that co-deliberative interaction is important and valuable. Further, one-on-one conversation provides a distinctive opportunity for both sides of the interaction. The deliberative duties are most at home here for that reason. B can clarify her view or amend it. The primary agent (A)¹ can question B, ensuring he understands her view, even seek her reasons for it. They each have deliberative duties that—when things go well—help to shape and motivate their conversational contributions. When the conversation concludes, and things have not gone well—B has chosen not to take advantage of this opportunity to fulfill her deliberative duties—B has opened herself up to complicity in A's actions.

¹ Recall from Chapter 1 that the *primary* agent is the wrongdoer who commits the primary wrong (the wrong in which the secondary wrongdoer is complicit). See footnote 14 in Chapter 1.

Why does it matter that B has an opportunity to clarify or amend her position? It might seem like the significance of the conversation is the possibility of mistake: outsiders to the conversation are more likely (than A) to misunderstand B's expressed stance. So it could be that the reason that the one-on-one conversation is particularly significant to whether or not B has a deliberative duty with respect to X's φ -ing depends on the likelihood of X's making a mistake about what B's adopted stance is. We might think that if C is in a position from which B's stance is more likely to be ambiguous, then B is not open (or is less open) to complicity in C's φ -ing because C should not have relied on (C's mistaken understanding of) B's stance. For example, if C is not party to the conversation, C may misunderstand the meaning of a statement of B's that he overhears.

Suppose instead of sharing with Brianna a cheating plan, Agatha discusses a study plan that involves reviewing sample exams. From the conversational context, Brianna knows that this does not constitute cheating: Agatha has mentioned earlier that her professor explicitly permitted reviewing sample exams as a study technique. Cynthia eavesdrops on their conversation but misses the earlier revelation that reviewing sample exams does not constitute cheating in this context. She mistakenly believes that Agatha plans to cheat and that Brianna has not objected. Cynthia takes both Agatha and Brianna to have *acted-as-though* cheating is permissible and goes on to cheat herself. If we agree that Agatha and Brianna have not failed in their deliberative duty with respect to Cynthia, we might think that this can be attributed to Cynthia's mistaken interpretation of Agatha's meaning (and thus also of Brianna's failure to correct).

But the possibility of mistake as the explanation of the conversation's significance seems at best incomplete for four main reasons. First and most importantly, it seems to get the order of explanation backwards: the possibility of mistake itself rests on something more fundamental, the fact that A and B (and not B and C) are involved in a conversation. In our example, Agatha and Brianna are talking with each other, and Cynthia is not part of that conversation. And it is this more

fundamental idea that is doing the work with respect to what B's deliberative duties are (to A and C) and to whether B is complicit in a particular wrongful action (of A's or C's). My second, third, and fourth worries are connected: it seems like what matters is whether or not B has a deliberative duty with respect to the primary wrongdoer and the relevant wrongdoing, and (2) sometimes if such a duty is in place A's (or C's) being mistaken about B's stance will be B's responsibility. That is, it will have been up to B to correct the mistake. If what mattered were the mistake (3) it seems we should think that there is complicity in the cases in which C happens to be *right* about B's stance and (4) it seems we should think there isn't complicity in the cases in which A happens to be *wrong* about B's stance (even if B knew A had misunderstood her and could have corrected her misunderstanding).

Rather than mistake, the conversational context is particularly significant because the fact that they are engaged in a dialogue makes it clear that A and B are engaged with one another on the relevant topic. They are speaking with each other and so also speaking *to* each other.² It does seem true that there is a greater chance that someone who is not party to the conversation (C) will be mistaken about B's expressed stance, but I think this greater chance of mistake is really just pointing to the more important difference between A and C—namely, that B was not talking to C. This does not mean that the deliberative duties *only* govern dialogue: it is still possible for B to fail in a deliberative obligation she has with respect to some possible action of C's, as we will see below. But what it does mean is that the core cases take place in dialogue.

Moving past dialogue, my thinking is a bit more tentative and so the arguments more speculative.

First, notice that because the deliberative view of complicity rests complicity on deliberative duty failure, the scope of deliberative complicity will depend on the scope of the deliberative duties. So it

² Moran also appeals to and defends the significance of one person talking with another (Moran 2018, 44–75). My arguments here are strongly influenced by his “assurance view,” according to which when B tells A something, B assumes responsibility for its truth by presenting herself as a kind of guarantor of its truth.

is to the deliberative duties we must look. Because deliberative duties contain both pieces—an agent and an action content—they are well-suited to answer these questions of scope. This does not mean that there will always be agreement about what the answer is, but it does mean that the disagreement will be traceable to disagreement about whether or not the potentially complicit agent failed in a deliberative duty she had with respect to *this* agent and *this* action of theirs. How should we understand their scope? I think the scope of the deliberative duties is both context and content sensitive, so I will say a bit about how and why.³

The context itself often seems to set the audience: there is a difference between Brianna and Agatha’s discussing Agatha’s potential cheating in their dorm room and having a loud conversation in the common room. When the conversation is more public, the participants should understand that their deliberative obligations extend more widely. Of course, we have already seen that Brianna’s duties to Agatha (the one *to whom* she is talking) are distinct from the duties she has toward those listening to their conversation, but it isn’t as though she has *no* deliberative obligations to those who might be listening when she and Brianna are speaking in a more public context. So to Cynthia, who eavesdrops on their private conversation, Agatha and Brianna owe very little in the way of a deliberative obligation.⁴ But to yet another party, Damon—who overhears their discussing cheating loudly in the common room—they owe a bit more (though still less than they owe each other).

As these examples help to bring out, the context of the conversation is significant to the scope of the deliberative obligations and so to the scope of moral complicity. Part of the context is the setting of the conversation, as the dorm room and common room examples help to illustrate.

The context also includes the relationship between the agents, such as their roles toward one

³ Here my thinking is influenced by Nissenbaum’s concept of “conceptual privacy” (2009, Chapter 7).

⁴ Could it be that they owe Cynthia nothing at all? Perhaps, but it doesn’t seem to be a hard and fast rule that nothing is ever owed to an eavesdropper. It seems to depend on (at least) the content of the relevant moral belief and how easy it was to determine that someone might be eavesdropping.

another and history of interaction, and other factors such as the culture and time period that help to set the salience of particular moral issues. In addition to the context, the content of the relevant moral belief can make a difference to the scope of the deliberative obligation. Some content is more public in nature and so lends itself to a deliberative obligation that extends to a wider audience; some is more private and so lends itself to a deliberative obligation that is narrower in scope. It also seems that the deliberative obligations might be more stringent in cases of content for which the moral truth is more easily accessible.

It also might matter if B knows that C attends to B's speech or activity: if B is aware that C is listening or paying attention, then B might have reason to be more careful and clearer in what she says and to pay more attention to the stance her action communicates. However, I don't think that B's knowledge in this area settles things. If B's conversation is in a context in which privacy was reasonable to expect but C is listening anyway, B seems to owe less clarity. If B happens to be unaware that C was listening but B is conducting a very loud conversation in a public space, B seems to owe more.

Together, the context and content of interaction help to set the scope of the deliberative obligation. They also help to answer some lingering questions. First, for how long does deliberative duty failure leave an agent open to complicity? For example, if Agatha were to cheat next year, would Brianna be complicit in that instance of cheating as well? Second, to which class of actions should we understand the failed deliberative duty to apply? This question is a version of the classic but vexing action description problem. One reason this matters is because if the appropriate description is quite specific, then the set of potential actions in which the secondary agent will be complicit will be quite small. But if the appropriate description is quite broad, then the set of potential actions will be rather large.

What resources does the deliberative conception of complicity bring to bear on these questions? We have noticed that the deliberative duties are most at home in conversation, where each agent can question the other, seeking clarification and seeking reasons. We could think of a failed deliberative duty like the spreading of a piece of gossip, sometimes analogized to squeezing the toothpaste out of tube: once it's out there is no way to get it back in. But this view seems too extreme because it does not seem to make space for changing one's mind or learning or growing morally. It seems that an agent who fails in her deliberative duty must have a way to close off the complicity to which she has opened herself. If she changes her mind about what she has said or expressed in the past, there will be ways to correct or cancel the stance she had previously adopted and clarify her new position. (Of course, not anything will be sufficient for cancelling—some action in addition to nice words will often be required.) And if we are far enough out from when she failed in her duty, even the cancelling may be unnecessary depending on the content of the relevant wrong. For example, it doesn't seem that one needs to call and correct what one said twenty years ago to a college acquaintance in order to avoid complicity in his wrongful actions. But if you and I are close, and I've changed my mind about something important, things seem different. On the other extreme, a view that terminated the complicity chain immediately, that only extended Brianna's complicity to Agatha's one instance of cheating, seems overly strong in the other direction. It also seems to miss the nuance and complexity of the deliberative duties.

When the deliberative view of complicity seeks an answer to whether B is complicit in some A's wrongful φ -ing, it begins by asking if B has failed in a deliberative duty she had with respect to that φ -ing. I see no reason to think that the temporal extension of one's liability to complicity will be the same in all cases. Rather, surely the temporal extension will depend on a number of factors, including the content of φ , the relationship between A and B (whether the relationship has continued or terminated, whether there is an implicit agreement not to discuss some matters because

each is already aware of the other's position, their disagreement, and their reasons for disagreeing), whether B has presented herself as having some expertise or special knowledge with respect to φ . Some cases will be clear, others hard. The account on its own will not answer all of our substantive questions about the difficult cases, but it does provide resources that bear on these questions. Sometimes it will be difficult to determine whether B failed in her deliberative duty with respect to A's φ -ing, but if these are also the cases in which it is difficult to determine whether B is complicit in A's φ -ing we will have reason to think that the account is on the right track. For example, it might sometimes be unclear whether the power differential in a relationship makes a difference to complicity, and in such cases it also seems unclear whether that same power differential makes a difference to the agent's deliberative duties. It can be hard to say whether a graduate student, for instance, has a deliberative duty to speak up against sexual harassment in her academic department, given her vulnerability and relative lack of power. But it is similarly difficult to say whether or not she is complicit. The idea here is that it is evidence in favor of the structural contention that moral complicity tracks deliberative duty failure if clarity of the cases rise and fall together.

Similar considerations bear on the description question. It is clear we cannot settle the question simply by direct appeal to the propositional content of what B says: a failure in one's deliberative duty could come about by implicature or indirect assertion. Nor can we rely on whatever B intended to convey, since B is responsible for some predictable false inferences on the part of her audience, even if she had no intention of causing them and was unaware of doing so.⁵ So we will have to look to the objective meaning of the stances expressed by B's conduct, whether through speech, action, or omission.

The action description question is a question of which actions of A's count as actions in which B has failed in her deliberative duty and so is complicit. We might imagine a set of core

⁵ For some helpful examples on this point (and on how complicated it is) see Shiffrin (2019).

actions of A's in which B would be complicit and some obvious possible actions of A's that are so far removed from the relevant deliberative duty that B would not be complicit in them.

Prospectively, we can consider all of the possible actions that A might perform as well as descriptions of those actions as a large possible set of action descriptions.⁶ In failing in her deliberative duty with respect to A, B has opened herself up to complicity in some subset of these possible actions, under such descriptions.⁷ The action description question asks how we determine the members of the subset. Let ψ be an arbitrary action of A's in this larger set, so it is a possible action of A's under a possible description. To decide whether ψ belongs to the subset that includes the actions in which B would be complicit if A performed them, we need to know if B has failed in her deliberative duty with respect to A's ψ -ing. So we have a principle of the form:

B is deliberatively complicit in A's φ -ing just in case B has failed in her deliberative duty with respect to A's φ -ing under that description.

Of course, this principle just pushes the question up a level. In order to determine whether B has opened herself up to complicity in a given instance of A's wrongful φ -ing, we need to know if B has failed in her deliberative duty with respect to A's φ -ing, and the principle does not even purport to answer *that* question. To answer that question, our inquiry takes us back to the arguments from Chapter 3 concerning the objective meaning of the stances B expresses through her conduct, including her speech, action, and omission. Since the next section turns from speech to action, I will return to this action description question below.

⁶ I prefer the prospective perspective because I think of B as aiming to fulfill her deliberative duty and thinking about how to do so. But you might also conduct the inquiry in a retrospective manner. Then the sets of A's actions will be smaller, including just A's actual actions and not all of A's possible actions.

⁷ I say "B has opened herself to complicity" since the set of actions described are A's *potential* actions, and B cannot be complicit in an action that A does not perform. So she will only be complicit in an even smaller set of actions than the described subset, namely, the set of actions within the subset that A goes on to perform.

4.5 Deliberative Duty Failure: Action

So far our focus has been on cases of speech. But if the arguments of Chapter 3 are right, then an agent's *action* can also constitute a failure in her deliberative duties. We have already begun to see that even though speech cases involve more speaker control over both audience and expressed content than action cases do, it can still be difficult to ascertain who counts as the relevant audience and what actions count among the scope when a potentially complicit agent fails in her deliberative duty. Both of these questions become even more difficult in the case of action, where the potentially complicit agent has considerably less control both over her audience and over the expressed content of her action. In the speech cases, I argued that what matters is whether B has failed in a deliberative duty she has with respect to A's φ -ing. How do we determine whether B has failed in such a deliberative duty due to B's actions?

Suppose Agatha and Brianna never have a conversation about cheating in which Agatha tells Brianna of her cheating plan. Instead, Brianna shows Agatha the easiest way to cheat, helping her to learn where to find old versions of exams for the class she is taking.⁸ Even though Brianna does not directly *say* that cheating is permissible, her actions clearly express her stance that cheating is permissible. Her actions cannot be explained in a way that is consistent with her having a commitment to the impermissibility of cheating. As we saw in Chapter 3, since her expression of the stance that cheating is permissible amounts to a failure in her deliberative duty, Brianna is complicit in Agatha's cheating. Part of what makes this case (and others similarly structured) clear is that as described only the two of them seem to be involved. Brianna's expressed stance seems unambiguous, if we consider what her action expresses *to Agatha*. But what about Cory, who sees Brianna help Agatha cheat and then goes on to cheat himself? Brianna's actions may not be so

⁸ Assume that this instance of reviewing old exams is cheating and is wrongful: the exams have been released without the consent of the instructor, and the instructor has made clear that reviewing them violates the course rules.

unambiguous to him, and so it is more difficult to tell if she has failed in her deliberative duty towards him and if she is complicit in his action.

Because features of our relationships make certain contents and audiences more salient, the relationship between A and B can make a difference to both of our scope questions (content and audience). Often, the deliberative duty is in place when there is some sort of hook or trigger, though this hook need not take the form of a conversation that has already begun or a question from A about whether to φ . A few more cases will help to illuminate the shape of the duty and the scope of the deliberative view of moral complicity.

Consider the case of a stranger, behind me in line outside the grocery store. We wait in line to enter, social distancing enforced as a result of the COVID-19 pandemic. He stands too close. He isn't wearing a mask. I ask him if he would consider putting one on. He refuses and shares his plan to do so once inside the store. I explain why that isn't sufficient, why he should be wearing one now, and would he please. He may still refuse or may relent, perhaps convinced or perhaps wanting me to leave him alone. By talking with him about the mask, I fulfill a deliberative duty.

This doesn't mean that not talking with him about mask wearing would certainly have constituted a failure of the duty. Since imperfect, the duty admits some discretion. It might be fulfilled in more than one way. Perhaps wearing my mask in line while outside fulfills my deliberative duty on its own because doing so expresses my stance that wearing a mask is important, is what's to be done under the circumstances. Similarly, *not* wearing a mask might constitute a failure in the deliberative duty, since not wearing a mask expresses one's stance that one need not wear a mask. If this is right, and if not wearing a mask constitutes a failure in the duty with respect to the person in line behind me in the grocery store, then if I fail to wear a mask, I will be deliberatively morally complicit if he goes on to not wear a mask the following day. I will be criticizable for his failure to mask appropriately.

Of course, substantive disagreement is possible here: I may be wrong that a failure to wear a mask constitutes a failure in one's deliberative duty with respect to the content of mask-wearing and the persons one encounters in the grocery store line. And I may be wrong about some of the other cases I will consider below. But as we have seen in previous chapters, such substantive disagreement is independent of my main structural contention that deliberative duty failure opens one up to moral complicity.⁹ The deliberative view of moral complicity gives us the resources to think about these sorts of cases and helps us to answer the question of moral complicity. To decide whether B is complicit in A's wrongful action, we need to decide whether B failed in a deliberative duty she had with respect to A and with respect to that wrongful action. A substantive disagreement over whether a particular deliberative duty is in play in a particular case will not pose a problem for the main view. An objection to the structural basis of my account would instead need to take a different form (perhaps defend a disconnect between deliberative duties and moral complicity).¹⁰ Such an objection might point to a case in which it seems there is deliberative duty failure but no complicity.

Return to the grocery store line. The fact that we were together in line for the grocery store made the importance of mask-wearing salient content between us. And the fact that such content was salient between us explains why I had a duty with respect to the person in line behind me and with respect to mask-wearing. Other content seems less salient or even not salient at all. It would be odd of me to turn around and begin to explain why lying is wrong. Or why the political practice of supporting schools via local property taxes is patently unjust. A deliberative duty that I bring

⁹ Scanlon makes a similar point about objections to his arguments concerning contractualism. An objector might disagree with Scanlon's evaluation of what people would reasonably reject, but this will not be an objection to the structural idea of contractualism but rather what Scanlon calls a "substantive disagreement" (2000, 186, 205, 238). I follow him in the use of this phrase.

¹⁰ Note that pointing to a case in which there is complicity but no deliberative duty failure will not work to form the basis of an objection of this kind, since the deliberative view of complicity is compatible with other forms of complicity explaining other cases.

these up and make them salient between us isn't in place just because we are close to each other in this grocery store line, and so I don't become complicit if he tells a lie or endorses unjust school funding practices.

Of course, there will be harder cases at the edges: is the fact that we're in the grocery store line sufficient to make salient the difficult moral questions surrounding the permissibility of meat eating? Probably not, unless some kind of hook or trigger gets it off the ground. We could imagine a few: the person behind me in line might ask how much meat I am buying for my freezer due to the impending meat shortages. Or he might observe me buying meat myself once we are inside the store. If eating meat is wrongful, I may be complicit in his meat-eating if I fail in my deliberative duty to discuss the reasons against meat-eating in the first instance and by expressing my stance that meat is to-be-bought in the second.

What do we learn from this sort of case? Content can become salient between two agents in a variety of ways, and then agents can fail in their deliberative duty with respect to that content once it has become salient. They can also *make* it salient between them, by bringing up the question for conversation. Sometimes, they can reject the opportunity to discuss the content, but this will not always be possible: engagement may be required, depending on the relationship and the content.

The content of mask-wearing became salient between the people in the grocery line due to the facts of the world on the occasion of them standing together in line. Even though they are strangers, there was a deliberative duty between them with respect to that particular content, due to the salience of that content between them.

We know that the deliberative duties do not require that we walk around telling all those we encounter all of the moral truths we know. This won't be effective, it won't be possible, and it isn't how practical co-deliberation is meant to go: it isn't how we treat those around us as agents, as persons who can act for the reasons they have. But we should try to help them get it right—try to

help them come to know the moral truth. The particular form this takes will vary based on our relationship with them and the content at issue.

4.6 Lingering Worries: Clean Hands and Others

Before concluding, I want to consider some lingering worries facing the proposed connection between deliberative duties and moral complicity. First, this proposal has the flavor of an unappealing sort of moral reasoning that focuses and worries only about keeping one's own hands clean and not about actually making a difference to evil and injustice in the world. In the case in which she isn't complicit, Brianna acts differently, but the results in the world are exactly the same as they would have been in the case in which I claim she *is* complicit. Why should we care that she isn't involved if the wrongful action takes place anyway? Or we might imagine Brianna asks us—why does her duty violation matter if the wrongful action would have taken place even if she had done what duty requires? Why follow the duty?

A second worry with this proposal is that it seems to undermine part of what makes moral complicity philosophically interesting in the first place. It seems at the beginning of the inquiry as though our intimate involvement with the lives of others can make us morally responsible for their wrongful actions even if we do nothing that is in itself wrongful. But an account that analyzes complicity in terms of duty fulfillment and failure runs the risk of failing to explore precisely what makes the topic of philosophical interest in the first place.

I think these worries are on to something: the best account of moral complicity should have the resources to distinguish itself from a my-hands-are-clean kind of moral reasoning. And it should also have something to say about when we should care about our involvement in the lives, actions, and choices of others. If all we end up with in the end is a directive to do our duties, focus only on

our own moral purity, and worry not at all about others, we will seem not to have gained much at all from thinking about complicity.

By focusing on the consequences, these worries get the order of explanation backward. How we decide *which* consequences are relevant is by looking to the duty, what it requires, and what it is about. Brianna is complicit in Agatha's cheating because of her duty failure: we explain the consequences that are attributable to her in terms of the duty failure itself. There isn't an independent account of which consequences she is responsible for.¹¹ Further, thinking about moral complicity in this way illustrates this interesting category of duties we might not have realized we have—duties that are focused on *others'* practical deliberation in a positive way, not just in way that prohibits lying and deception. And these deliberative duties are fundamentally about our engagement with others and with the practical deliberation of others. So, the second worry misses its target in part because it is in failing in the deliberative duties that we get interaction with others wrong. To focus our inquiry into complicity on the deliberative duties is not to fail to think about our entanglements with others. On the contrary. The deliberative duties explain what is required when we are involved with the deliberation and actions of others, and so they do engage with the part of the concept of moral complicity that makes it philosophically interesting and challenging and provide a compelling explanation.

The deliberative duties govern our interactions with other agents and their moral reasoning. And so, they are intimately connected with other agents' practical reasoning and with their actions. This makes the deliberative duties uniquely well-positioned to explain complicity, which is also intimately connected with other agents' reasoning and action. Other sorts of duties will be at a significant disadvantage in playing this kind of explanatory role. While failure in non-deliberative duties might open an agent up to responsibility for consequences in a similarly asymmetrical style,

¹¹ Herman makes a similar point in her response to Callard on anger (2020).

they do not primarily govern our interactions with the other agents' moral reasoning.¹² For example, consider the duty of promise-keeping. Suppose Brianna fails to keep her promise to Agatha, perhaps not picking her up from the airport despite having promised that she would. It seems such duty failure might open Brianna up to responsibility for some of the consequences that follow from her failure, including perhaps an obligation to bear some or all of the cost of Agatha's air transfer. But suppose that Agatha lies to her cab driver in an attempt to get a reduced fare, claiming to be someone she isn't. Brianna would then not be complicit in Agatha's wrongful action, despite having broken a promise she made to Agatha. Why not? It seems to be because she does not bear the right sort of connection to Agatha's lying. So a duty like promise-keeping, though of course connected with other agents because interpersonal (in the usual case, setting aside promises made to the self), does not seem well-suited to explain moral complicity. If, on the other hand, Brianna had told Agatha that lying in such cases was perfectly fine—that trying to take advantage of one's ability to deceive others in these kinds of circumstances was even good or to be valued—we should understand her as complicit because of the intimate connection between her duty failure and the relevant wrongful action. (And we might again note that neither causal contribution nor intentional participation seems well-situated to sort the cases.)

In sum, this section has offered one distinctive reason for thinking that the deliberative duties are well-suited to play the role of determining and explaining moral complicity—their intimate connection with the reasoning and actions of another agents.

¹² An important caveat here. If the arguments of Chapter 3 are right, then failures in non-deliberative duties will *also* constitute a failure in at least some of one's deliberative duties, since they will be instances of *acting-as-though* such failure is permissible. So an agent *can* become complicit by failing in a non-deliberative duty not because the failing in such a duty makes her complicit but because it *also* amounts to her failure in her *deliberative* duty. Having adopted and communicated the incorrect stance toward the duty, she has opened herself up to complicity for another's failure in that duty (not just any other's failure, though—more will need to be said about who counts as her audience). In this discussion in the body, I set aside this complication, since I discuss the case of action in §4.5.

4.7 Conclusion

Philosophical accounts of moral complicity owe some answers to our questions of scope. Over what set of activities does our complicity range? Intentional participation views fare well along this dimension, and causation-based views and deliberative views provide less effective and compelling answers. In this chapter, I have articulated some of the resources the deliberative view can offer to resolve these questions of scope, but some such questions remain unresolved. In the end, I am not convinced we should expect sharp distinctions: it may be that sometimes it is just very hard to know if a deliberative duty is in place and thus whether one is complicit in another's wrongful action. It may be that sometimes the best course of action in a case of uncertainty is to err on the side of caution—walk away, clarify one's position, avoid potentially ambiguous behavior. In any event, the deliberative view offers new resources that bear on our understanding of the scope of our complicity: if we are interested to know the actions over which our complicity ranges, we should look to the scope of our deliberative duties.

Conclusion

In this dissertation, I have defended a kind of moral complicity that is fundamentally deliberative in nature. It does not depend on causal contribution nor intentional participation. So an agent might be morally complicit in the wrongdoing of another agent even if she does not intend to participate in that wrongdoing and does not increase its likelihood. When Cindy complains that you are complicit in Albert's breaking a promise to her, your defenses that you didn't intend his breach nor increase the chances he would breach fall flat. You are morally complicit in his promise-breaking in virtue of telling him that there is nothing at all wrong with breaking a promise because to do so was to fail in a deliberative duty you had to help him think and act rightly.

Deliberative complicity is not limited to cases of speech. Our actions can also constitute deliberative duty failures and so open us up to moral complicity. When Betsy drives Adam to the bank knowing he will rob it, she *acts-as-though* bank-robbing is permissible. Even if she tells Adam he shouldn't rob, she is complicit in his robbing in virtue of the deliberative duty failure constituted by her driving him to the bank.

Understanding our complicity as deliberative instead of causal will not answer all of our questions about our moral responsibility for others' wrongdoing, since we will still need to know when our action, speech, or silence constitutes deliberative duty failure. For instance, we may still wonder whether simple participation in a global economy is sufficient for deliberative duty failure with respect to injustices committed against the global poor. In future work, I hope to think more about deliberative duties in particular contexts, since I think the context of the interaction and the content of relevant moral belief will both play a role in the scope of the deliberative duty and thus the scope of deliberative moral complicity.

And yet a shift to an understanding of complicity as deliberative rather than causal *will* offer helpful resources as we engage with these questions. In particular, a deliberative conception of

moral complicity and the default entitlement defended in Chapter 2 can help us to see how we might balance two seemingly-conflicting values—respect for the autonomy and agency of others and avoiding complicity in their wrongdoing.

A deliberative account of moral complicity can help us to reimagine our understanding of our moral impact on the world and so help us make progress on difficult questions of our complicity in an increasingly global society. By helping others to see the moral truth, I can not only avoid complicity in wrongdoing but also support others who aiming act rightly. Through my speech and action, I express my stance on moral issues, and ensuring that those stances are both correct and in line with my actual thinking is important not only to my own moral flourishing but also to the moral success of those around me.

Bibliography

- Anderson, Elizabeth S., and Richard H. Pildes. 2000. "Expressive Theories of Law: A General Restatement." *University of Pennsylvania Law Review* 148: 1503–76.
- Anscombe, G. E. M. 2000. *Intention*. 2nd ed. edition. Cambridge, Mass.: Harvard University Press.
- Baron, Marcia. 2003. "Manipulativeness." *Proceedings and Addresses of the American Philosophical Association* 77 (2): 37–54. <https://doi.org/10.2307/3219740>.
- Basu, Rima. 2019. "What We Epistemically Owe to Each Other." *Philosophical Studies* 176 (4): 915–31. <https://doi.org/10.1007/s11098-018-1219-z>.
- Berenstain, Nora. 2016. "Epistemic Exploitation." *Ergo: An Open Access Journal of Philosophy* 3: 569–590. <https://doi.org/10.3998/ergo.12405314.0003.022>.
- Bratman, Michael. 1987. *Intention, Plans, and Practical Reason*. David Hume Series. Cambridge, Mass.: Harvard University Press.
- Burge, Tyler. 2003. "Perceptual Entitlement." *Philosophy and Phenomenological Research* 67 (3): 503–48. <https://doi.org/10.1111/j.1933-1592.2003.tb00307.x>.
- . 2020. "Entitlement: The Basis for Empirical Epistemic." In *Epistemic Entitlement*, edited by Peter J. Graham and Nikolaj J. L. L. Pedersen, 37–142. Oxford: Oxford University Press. <https://doi.org/10.1093/oso/9780198713524.003.0002>.
- Burwell v. Hobby Lobby. 2014 1. U.S. Supreme Court.
- Callard, Agnes. 2020. "The Philosophy of Anger." In *On Anger*, edited by Agnes Callard. Boston Review. <http://bostonreview.net/forum/agnes-callard-philosophy-anger>.
- Carpiniello, Mia. 2001. "Striking a Sincere Balance: A Reasonable Black Person Standard for 'Location Plus Evasion' Terry Stops." *Michigan Journal of Race and Law* 6 (2): 355–87.
- Davidson, Donald. 1963. "Actions, Reasons, and Causes." *The Journal of Philosophy* 60 (23): 685. <https://doi.org/10.2307/2023177>.
- Dotson, Kristie. 2018. "On the Way to Decolonization in a Settler Colony: Re-Introducing Black Feminist Identity Politics." *AlterNative: An International Journal of Indigenous Peoples* 14 (3): 190–99. <https://doi.org/10.1177/1177180118783301>.
- Frowe, Helen, and Jonathan Parry. 2019. "Wrongful Observation." *Philosophy & Public Affairs* 47 (1): 104–37. <https://doi.org/10.1111/papa.12142>.
- Gardner, John. 2004. "Christopher Kutz, Complicity: Ethics and Law for a Collective Age." *Ethics* 114 (4): 827–30. <https://doi.org/10.1086/420729>.

- . 2007. “Complicity and Causality.” *Criminal Law and Philosophy* 1 (2): 127–41.
<https://doi.org/10.1007/s11572-006-9018-6>.
- Grice, Paul. 1991. “Logic and Conversation.” In *Studies in the Way of Words*, 22–40. Cambridge, Mass.: Harvard University Press.
- Hawkins, Derek, and Morgan Krakow. 2019. “Walmart, Site of Recent Shootings, Has a Complicated History with Guns.” *The Washington Post*, August 4, 2019.
- Helmreich, Jeffrey S. 2015. “The Apologetic Stance.” *Philosophy & Public Affairs* 43 (2): 75–108.
<https://doi.org/10.1111/papa.12053>.
- Herman, Barbara. 2011. “A Mismatch of Methods.” In *On What Matters, Vol. II*, edited by Samuel Scheffler, 1 edition, 83–115. New York: Oxford University Press.
- . 2019. “Being Prepared: From Duties to Motives.” In *Oxford Studies in Normative Ethics Volume 9*, edited by Mark Timmons, 9–28. Oxford Studies in Normative Ethics. New York: Oxford University Press.
- . 2020. “What’s Past Is Prologue.” In *On Anger*, edited by Agnes Callard. Boston Review.
- Hieronymi, Pamela. 2008. “The Reasons of Trust.” *Australasian Journal of Philosophy* 86 (2): 213–36.
<https://doi.org/10.1080/00048400801886496>.
- Holmes, Oliver Wendell. 1909. *The Common Law*. Little, Brown.
<http://archive.org/details/commonlaw00holmgoog>.
- Honan, Edith. 2019. “Walmart Says It Sells 20% of Ammunition in US, Defends Gun Sales after Mass Shootings,” August 15, 2019. <https://abcnews.go.com/US/walmart-defends-gun-sales-shootings-stores/story?id=64993374>.
- Jackson, Frank. 2014. “Procrastinate Revisited.” *Pacific Philosophical Quarterly* 95 (4): 634–47.
<https://doi.org/10.1111/papq.12051>.
- Jackson, Frank, and Robert Pargetter. 1986. “Oughts, Options, and Actualism.” *The Philosophical Review* 95 (2): 233–55. <https://doi.org/10.2307/2185591>.
- Kant, Immanuel. 1999. *Practical Philosophy*. Edited by Mary J. Gregor. New Ed edition. Cambridge: Cambridge University Press.
- Korsgaard, Christine M. 1986. “The Right to Lie: Kant on Dealing with Evil.” *Philosophy & Public Affairs* 15 (4): 325–49.
- Kutz, Christopher. 2000. *Complicity: Ethics and Law for a Collective Age*. Cambridge Studies in Philosophy and Law. New York: Cambridge University Press.
- . 2007. “Causeless Complicity.” *Criminal Law and Philosophy* 1 (3): 289–305.

- <https://doi.org/10.1007/s11572-006-9026-6>.
- Lackey, Jennifer. 2020. "The Duty to Object." *Philosophy and Phenomenological Research* 101 (1): 35–60.
<https://doi.org/10.1111/phpr.12563>.
- Lepora, Chiara, and Robert E. Goodin. 2015. *On Complicity and Compromise*. Reprint edition. New York: Oxford University Press.
- Manne, Kate. 2014. "Non-Machiavellian Manipulation and the Opacity of Motive." In *Manipulation: Theory and Practice*, edited by Christian Coons and Michael Weber, 221–45. New York: Oxford University Press.
- May, Larry. 2010. "Complicity and the Rwandan Genocide." *Res Publica* 16 (2): 135–52.
<https://doi.org/10.1007/s11158-010-9112-4>.
- McMillon, Doug. 2019. "McMillon to Associates: Our Next Steps in Response to the Tragedies in El Paso and Southaven."
<https://corporate.walmart.com/newsroom/2019/09/03/mcmillon-to-associates-our-next-steps-in-response-to-the-tragedies-in-el-paso-and-southaven>.
- Moran, Richard. 2018. *The Exchange of Words: Speech, Testimony, and Intersubjectivity*. New York: Oxford University Press.
- Murphy, Liam B. 2003. *Moral Demands in Nonideal Theory*. Oxford: Oxford University Press.
- Nagel, Thomas. 1979. "Moral Luck." In *Mortal Questions*, 24–38. Cambridge [Eng.]: Cambridge University Press.
- Nissenbaum, Helen. 2009. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. 1st edition. Stanford, Calif: Stanford Law Books.
- O’Neil, Collin. 2012. "Lying, Trust, and Gratitude." *Philosophy & Public Affairs* 40 (4): 301–33.
<https://doi.org/10.1111/papa.12003>.
- Pallikkathayil, Japa. 2011. "The Possibility of Choice: Three Accounts of the Problem with Coercion." *Philosophers’ Imprint* 11.
- Peterson, Linda L. 1999. "The Reasonableness of the Reasonable Woman Standard." *Public Affairs Quarterly* 13 (2): 141–58.
- Petersson, Björn. 2013. "Co-Responsibility and Causal Involvement." *Philosophia* 41 (3): 847–66.
<https://doi.org/10.1007/s11406-013-9413-x>.
- Sartorio, Carolina. 2007. "Causation and Responsibility." *Philosophy Compass* 2 (5): 749–65.
<https://doi.org/10.1111/j.1747-9991.2007.00097.x>.

- Scanlon, T. M. 2000. *What We Owe to Each Other*. Revised edition. Cambridge, Mass.: Belknap Press: An Imprint of Harvard University Press.
- . 2008. *Moral Dimensions: Permissibility, Meaning, Blame*. Cambridge, Mass.: Belknap Press of Harvard University Press.
- Shiffrin, Seana Valentine. 2014. *Speech Matters: On Lying, Morality, and the Law*. Carl G. Hempel Lecture Series. Princeton: Princeton University Press.
- . 2019. “I—Learning about Deception from Lawyers.” *Aristotelian Society Supplementary Volume* 93 (1): 69–90. <https://doi.org/10.1093/arisup/akz010>.
- Smart, J. J. C., and Bernard Williams. 1973. *Utilitarianism: For and Against*. Cambridge [Eng.]: University Press.
- Stilz, Anna. 2011. “Collective Responsibility and the State.” *Journal of Political Philosophy* 19 (2): 190–208. <https://doi.org/10.1111/j.1467-9760.2010.00360.x>.
- Susskind, Randall S. 1993. “Race, Reasonable Articulate Suspicion, and Seizure.” *American Criminal Law Review* 31 (2): 327–50.
- “The Principles Governing Cooperation and Catholic Health Care: An Overview.” 2013. *Health Care Ethics USA* 21 (1). <https://www.chausa.org/docs/default-source/general-files/a7ca2015424c4be8b14e99825058c8931-pdf.pdf?sfvrsn=0>.
- Williams, Bernard. 2002. *Truth & Truthfulness: An Essay in Genealogy*. Princeton: Princeton University Press.
- Wood, Allen W. 2008. *Kantian Ethics*. Cambridge ; New York: Cambridge University Press.
- Wrenn, Chase B. 2007. “Why There Are No Epistemic Duties.” *Dialogue: Canadian Philosophical Review / Revue Canadienne de Philosophie* 46 (1): 115–36. <https://doi.org/10.1017/S001221730000158X>.

ProQuest Number: 28543335

INFORMATION TO ALL USERS

The quality and completeness of this reproduction is dependent on the quality and completeness of the copy made available to ProQuest.



Distributed by ProQuest LLC (2021).

Copyright of the Dissertation is held by the Author unless otherwise noted.

This work may be used in accordance with the terms of the Creative Commons license or other rights statement, as indicated in the copyright statement or in the metadata associated with this work. Unless otherwise specified in the copyright statement or the metadata, all rights are reserved by the copyright holder.

This work is protected against unauthorized copying under Title 17, United States Code and other applicable copyright laws.

Microform Edition where available © ProQuest LLC. No reproduction or digitization of the Microform Edition is authorized without permission of ProQuest LLC.

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 - 1346 USA