

Forthcoming in *The Journal of Value Inquiry*

The Just Soul

Jeff Sebo
New York University

Abstract

Many philosophers think that, if your day self and night self are physically, psychologically, and narratively continuous with each other, then they are the same basic unit of moral concern. But I argue that your day self and night self can share all of these relations and still be different units of moral concern, on the grounds that they can share all of these relations and still be in the circumstances of justice with each other. I then argue that this conception of the scope of morality has revisionary, but ultimately plausible, implications for the morality of self-binding. For example, it implies that your day self and night self have a prima facie moral duty not to coerce or physically restrain each other in order to get what they want. But it also implies that they are morally permitted to coerce and physically restrain each other much more often, and with respect to many more issues, than, say, you and your friend are.

1. Introduction

Consider the following two cases:

1. You and your friend Doug are on a road trip, and you want to get an early start in the morning. But you know that this will be impossible if Doug drinks tonight, and you also know that, when night rolls around, Doug will want to drink. What should you do? You have two options. First, you can compromise and coordinate with Doug: for example, you can promise to stay out late tomorrow night if he gets to bed early tonight. Second, you can coerce and physically restrain Doug. For example, you can pour all his beer down the drain, slip him a sleeping pill, lock him in his room overnight, and threaten to do much more in the future if he even *thinks* about drinking tonight.
2. You are on a road trip, and you want to get an early start in the morning. But you know that this will be impossible if you drink tonight, and you also know that, when night rolls around, you will want to drink. What should you do? You have two options. First, you can compromise and coordinate with yourself: for example, you can plan to stay out late tomorrow night if you get to bed early tonight. Second, you can coerce and physically restrain yourself: for example, you can pour all your beer down the drain, take a sleeping pill, lock yourself in your room overnight, and plan to do much more in the future if you even *think* about drinking tonight.

What does morality have to say about these cases? Most philosophers would evaluate them differently. In the first case, they would say that, since you and Doug are different units of moral concern (i.e. different bearers of moral duties and rights), you have a moral duty to compromise and coordinate with each other rather than coerce or physically restrain each other. In the second case, however, most philosophers would say that, since your day self and night self are the same unit of moral concern, they are morally permitted to *either* compromise and coordinate with each other *or* coerce and physically restrain each other. And when asked why your day self and night self are the same unit of moral concern whereas you and Doug are not, most philosophers would reply in one of two ways. On one hand, philosophers who think that personal identity is what matters in morality (i.e., that being the same person is necessary and sufficient for being the same unit of moral concern) would reply by saying that your day self and night self are the same person, whereas you and Doug are not. On the other hand, philosophers who think that personal identity is *not* what matters in morality would likely reply by saying that your day self and night self are physically, psychologically, and/or narratively continuous with each other, whereas you and Doug are not.ⁱ

I will not consider the first kind of reply here: I agree with Derek Parfit (and many other philosophers) that personal identity is not what matters in morality, and I will be taking that assumption for granted in what follows.ⁱⁱ My aim in this paper, then, is to argue that, *if* we accept that personal identity is not what matters in morality, then we should *also* accept that physical, psychological, and narrative continuity are not what matters in morality either, since, as with personal identity, (a) *x* and *y* can share all of these relations and still be in *the circumstances of justice* with each other (i.e., the circumstances that make justice both possible and necessary to try to achieve),ⁱⁱⁱ (b) if *x* and *y* are in the circumstances of justice with each other, then they have

duties of justice to each other (i.e., they have a duty to jointly commit to a fair set of laws and to bind each other only insofar as they *have* to, in order to enforce those laws), and (c) if x and y have duties of justice to each other, then they are, by definition, different units of moral concern.

This conception of the scope of morality will have revisionary implications for the morality of self-binding. Specifically, if being in the circumstances of justice is sufficient for being different units of moral concern, then your day self and night self in this case are different units of moral concern, with a *prima facie* moral duty not to coerce or physically restrain each other, after all. I will argue, however, that this implication is not as revisionary as it might first appear, for the following reason: even if your day self and night self are different units of moral concern, they still have a much more intimate relationship with each other than you and Doug have (precisely because they are physically, psychologically, and narratively continuous with each other whereas you and Doug are not), and therefore they are still morally permitted to coerce and physically restrain each other much more often, and with respect to many more issues, than you and Doug are in practice. This discussion will underscore a further similarity between personal identity, on one hand, and physical, psychological, and narrative continuity, on the other hand: even if these relations are not sufficient to place our interactions with our past and future selves beyond the scope of morality, they are still sufficient to make our interactions with, and therefore our moral duties to, our past and future selves very different than they would otherwise be, in the same kind of way that the relation 'being in the same family as' is sufficient to make our interactions with, and therefore our moral duties to, other people very different than they would otherwise be. The upshot of this discussion, then, will be that we are right to think that these relations are morally significant but wrong to think that they are, in and of themselves, what matters in morality.

I will proceed as follows. In sections 2-5, I will present a thought experiment in which you and Doug gradually become a day self and a night self in the same person, and I will argue that, intuitively, you still count as different units of moral concern even when this process is complete. Then, in section 6, I will argue that this thought experiment supports, and finds support in, the view that, if x and y are in the circumstances of justice with each other, then they are different units of moral concern whether or not they are physically, psychologically, or narratively continuous with each other.

Before I begin, I should make a few clarifications about my terminology and strategy in what follows. First, even though I will not be considering the possibility that personal identity is what matters in morality here, it will still be useful for me to be able to talk about persons. So, I will assume that a standard theory of personal identity is true. That is, I will assume that physical, psychological, and narrative continuity are jointly sufficient for personal identity and, therefore, that your day self and night self in this case are, in fact, the same person. I want to emphasize, however, that nothing of substance will depend on this assumption. So if you think that physical, psychological, and narrative continuity are *not* jointly sufficient for personal identity, then you can translate all of my claims about *persons* into claims about *humans* without affecting the substance of my argument at all. For example, when I say that your day self and night self can be different units of moral concern despite being the same person, you can read me as saying that your day self and night self can be different units of moral concern despite being the same human, if you prefer.

Second, and relatedly, I will sometimes describe persons as having temporal parts, and I will also sometimes describe some of these temporal parts as selves.^{iv} For example, as we have seen, I will use ‘your day self’ to refer to you during the day and ‘your night self’ to refer to you

during the night. But as with my assumption about personal identity, nothing of substance will depend on this terminological choice. So if you think that persons endure rather than perdure, then you can translate all of my claims about your day self and night self into claims about, say, you during the day and you during the night without affecting the substance of my argument at all. For example, when I say that your day self has a prima facie moral duty not to coerce or physically restrain your night self (and vice versa), you can read me as saying that you-during-the-day have a prima facie moral duty not to coerce or physically restrain yourself-during-the-night (and vice versa), if you prefer.

Finally, since my strategy is to take you and Doug through a series of changes and argue that, intuitively, none of these changes is sufficient to take away your duties of justice to each other, my argument might appear to trade essentially on a sorites case. However, this appearance would be misleading. The idea is not to take you through a series of *small* changes and ask, each step of the way, whether a change *this small* could possibly take away your duties of justice to each other. (Indeed, each of these changes will be very big.) Rather, the idea is to isolate each of the main differences between interpersonal relationships, on one hand, and intrapersonal relationships, on the other hand, and to argue that none of these differences is morally relevant for present purposes. You and Doug have duties of justice to each other, I will argue, *not* because you have separate bodies and brains, chains of memory, streams of psychological connectedness, or self-narratives, but *rather* because you are in the circumstances of justice with each other; and this remains true even when you are a day self and night self in the same person.

With these clarificatory points out of the way (but keeping them in mind as we proceed), we can now begin the thought experiment.

2. Sharing a body and brain

Suppose that you and Doug have completely different conceptions of the good life. On one hand, you like to wake up early and start the day with a light breakfast and strong cup of coffee. Then you get to work. You write poetry for a living, so you spend most of your day reading and writing in your study. Then you like to relax: you cook dinner and watch TV or read a book. The one exception is: every now and then you like to invite some friends over for dinner. You eat good food, drink good wine, and have good conversation. Otherwise, though, you rarely drink and you always get to bed early, because you hate feeling sluggish and hung over in the morning; you like to be able to jump out of bed and do it all again.

On the other hand, Doug likes to sleep in and then start the day with a six-pack of beer. Then he gets to work. He plays punk rock for a living, so he spends most of the night writing, rehearsing, and recording in his home studio. Then he likes to have fun: he goes out to bars and clubs and shows, and then he wanders around the city and talks to strangers until five or six in the morning. The one exception is: every now and then, he likes to invite a bunch of people over for a party – not a dinner party, but a real party. He puts on some music, and they all drink and dance and sing until they pass out, or go to the diner for some breakfast.

Now imagine that one day, an evil scientist kidnaps you and Doug, takes complete brain scans of both of you, and destroys your bodies and brains. He then builds a new body and gives this body multiple personalities. Specifically, the evil scientist makes it the case that, from 6am to 6pm everyday, this new body is psychologically continuous with *you*: it has all of your memories, beliefs, desires, aims, intentions, habits, and so on. And from 6pm to 6am everyday, this new body is psychologically continuous with *Doug*, in all the same ways. This means that

this new body's two personalities are completely *psychologically discontinuous* with each other, so that each "blacks out" when the other takes over. When the clock strikes six in the morning, you wake up with no memory at all of anything that Doug thought, felt, or did last night. And when the clock strikes six at night, Doug wakes up with no memory at all of anything that *you* thought, felt, and did during the day, in the same kind of way.^v

Practically speaking, how does sharing a body and brain affect your relationship with Doug? In the same way that sharing a city, a house, a room, or any other piece of territory would. It makes your lives much more intertwined, in a way that creates new opportunities for conflict and disagreement. Before it was easy for you and Doug to live and let live: you could live one way and Doug could live another way, and there was no real conflict between you. But now that you share a body and brain, all of this has changed. If Doug stays up all night, it will be hard for you to wake up early and have a productive day. If he drinks and smokes and eats fast food all the time, it will be hard for you to maintain your high standards of health. If he gets tattoos and piercings all over his face, it will be hard for you to present yourself as a serious professional to your colleagues. And so on. Similarly: if you stay up all day, then it will be hard for Doug to wake up late and have a productive night. If you never let him drink or smoke or eat fast food at all, it will be hard for him to have fun. If you insist on keeping your body tattoo and piercing free, it will be hard for him to present himself as a counter-cultural punk to his friends. And so on. The upshot is that, since your body and brain play such an important role in both of your lives, and since you each have such different preferences about how to treat your body and brain, your *sharing* a body and brain makes it much harder for both of you to live exactly the way that you want.

What should you do about this? You have two options. First, you can try to live and let live. You can talk through your problems (maybe by passing notes back and forth) and jointly commit to a fair set of “bodily laws”. For example, maybe you can agree to sleep for four hours a day if Doug does the same. You can agree to lower your standards of health a bit if Doug raises his a lot. You can agree to let Doug decorate *parts* of your body with tattoos and piercings, provided that you can cover them up or take them off when you need to. And so on.

Alternatively, you can try to live the way that you want by any means necessary. For example, if Doug wants to be able to drink, smoke, and eat fast food each night without having to give up *anything* in exchange, then he can take the following approach: he can destroy all your possessions and get tattoos and piercings all over your face. Then he can stay up all night drinking, smoking, eating fast food, and doing a cocktail of drugs that will make you catatonic the next day, if you even wake up at all. Then he can write you a note that says: "If you even *think* of doing anything other than *exactly* what I tell you to in the future, I will personally see to it that you never write another book for the rest of your life."

The question is: What does morality have to say about this? Is Doug morally permitted to coerce and physically restrain you in these ways on the grounds that you share a body and brain with him now? Intuitively, the answer is clearly no. You may be different *personalities* in the same *human being* rather than, say, different *human beings* in the same *house*, but you are nevertheless in the same kind of predicament as two human beings in the same house: you disagree about how to live and yet you also have to live *together*. Thus, intuitively, you have a duty to live and let live. That said, sharing a body and brain *does* change your moral relationship in at least one important way: you now have to compromise and coordinate much more often, and with respect to many more issues, in order to live and let live. And this means that, first, you

have a *duty* to compromise and coordinate much more often and with respect to many more issues, and, second, you will probably have a right to *bind* each other much more often and with respect to many more issues. For example, if a fair compromise between you and Doug requires you each to sleep for four hours a day, then you have a right to ask Doug to get to bed by 2am, and, in extreme cases, to *make* him get to bed by 2am, for example by putting a sleeping pill in his food; and vice versa. This is of course a recent development: before, it would have been invasive (and weird) for you to interfere with each other in these ways. But now that you share a body and brain, you can interfere with each other in these ways without being invasive – not because you can do whatever you want to each other now, but rather because you may *have* to interfere with each other in these ways in order to protect your own rights.

3. Sharing a chain of memory

Now imagine that the evil scientist returns. He kidnaps you and Doug and rewires your brain so that you now share a chain of memory. This means that the black outs stop: when the clock strikes six in the morning, you wake up and remember everything that Doug thought, felt, and did last night as though *you* thought, felt, and did it. And when the clock strikes six at night, Doug wakes up and has the same kind of experience.^{vi} But imagine that your psychological connectedness stops there: you still have all the same beliefs, desires, and habits that you had before, and so does Doug.

Practically speaking, how does sharing a chain of memory affect your relationship with Doug? In two main ways. First, it affects your relationship in the same kind of way that sharing a body and brain did: it makes your lives even more intertwined, in a way that creates new

opportunities for conflict and disagreement. After all, before this step in your fusion, you may have had to compromise and coordinate with respect to how to treat your *body and brain*, but at least you each still had full autonomy with respect to your own *mind*, i.e., at least you had the option of keeping your *thoughts and feelings* to yourselves if you wanted to. But now that you share a chain of memory, all of this has changed. For example, you and Doug might disagree about whether he is permitted to tell everyone all your deepest, darkest secrets – secrets that you never would have told him about previously but which you now have no choice but to share with him. Similarly, you might disagree about whether Doug is permitted to perform actions the memory of which you find especially troubling and distracting – actions that you never would have asked him about previously but which you now have no choice but to learn about. And the upshot is the same as before: since your thoughts and feelings play such an important role in both of your lives, and since you each have different preferences about how to treat your thoughts and feelings, your sharing a chain of memory makes it even harder for both of you to live exactly the way that you want. And, insofar as you and Doug disagree about what to do, you have the same two options as before: you can either try to live and let live, or you can try to live the way that you want by any means necessary.

Second, however, sharing a chain of memory also affects your relationship with Doug in the same kind of way that the next two steps in your fusion will: it makes it easier for you to live and let live. How? By allowing you to use different *methods* for interacting with each other. Specifically, before this step in your fusion you had to use “third-personal” methods for interacting. For example, if you wanted to compromise and coordinate with each other, then you had to do all of the following: First you had to think about what would make for a fair compromise between you. Then you had to put it all into language. Then you also had to try to

interpret everything each other said. But now that you and Doug share a chain of memory, you can dispense with all this. For example, if you want to compromise and coordinate with each other, then all you have to do is think about what would make for a fair compromise between you and then form an intention to act on this compromise. For example, you can intend to have a couple of drinks tonight and then come home at a reasonable hour with the expectation that Doug will remember your intention and act on it (perhaps with some minor revisions) – not because he *prefers* to do this (he prefers to stay out much later and drink much more), but rather because he recognizes that you formed this intention in the spirit of compromise, and he also recognizes the importance of your being able to think and act on behalf of the human being as a whole, given how many issues you have to compromise and coordinate with respect to every day.

But even if your chain of memory makes compromising and coordinating much easier (and also makes coercing and physically restraining much harder, since you have no way to keep secrets from each other anymore), you can still try to coerce or physically restrain each other if you want to. For example, if you want to be able to make Doug wake up early each morning without having to give up anything in exchange, then you can take the following approach: you can destroy all his possessions, get all his tattoos removed, and take all his piercings out of your body. Then you can write to everyone he knows and tell them all his deepest, darkest secrets. Then you can take a sleeping pill that will make him catatonic that night, if he even wakes up at all. Finally, you can tell Doug, merely by framing the intention in your own mind: “This is *my* body and *my* mind. If you even *think* about doing anything other than *exactly* what I tell you, I will personally see to it that you never write another song for the rest of your life.” And you can keep doing this, day in and day out, until Doug finally relents, miserable but out of options.

What does morality have to say about this? Are you morally permitted to coerce and physically restrain Doug in these ways on the grounds that you share a chain of memory now? Intuitively, the answer is still no. You may know everything that each other thinks, feels, and does now, and you may also have a new, first-personal way of talking to each other as a result. But nothing else has changed: you still disagree about how to live (maybe even more than before), and yet you still have to live *together*. Thus, intuitively, you still have a duty to live and let live. Indeed, if anything, sharing a chain of memory seems to have changed your moral relationship in the same kind of way that sharing a house and a body and brain did: you have to compromise much more often, and with respect to many more issues, in order to live and let live; and therefore you have a *duty* to compromise much more often and with respect to many more issues, as well as a right to *bind* each other much more often and with respect to many more issues. For example, if a fair compromise between you and Doug requires you each not to perform any actions the memory of which the other finds especially troubling, then you now have a right to ask Doug to not, say, eat lima beans, and in extreme cases to *make* him not eat lima beans; and vice versa. Fortunately, as I said before, these negotiations will be much easier to carry out now that you and Doug can talk to each other without ever opening your mouth or picking up a pen. But they will also have to cover pretty much every aspect of your lives now, with the result that, morally speaking, almost *nothing* is up to you alone anymore.

4. Sharing a high degree of psychological connectedness

Fourth, imagine that the evil scientist kidnaps you and Doug again and rewires your brain so that you now share a high degree of psychological connectedness. This means that, if you form a

particular belief, desire, or aim, then Doug is very likely to *inherit* this particular psychological state, and vice versa. For example, when the clock strikes six in the morning, you wake up and not only *remember* everything that Doug thought, felt, and did over the past twelve hours but also *endorse* many more of these thoughts, feelings, and actions than you did before. And when the clock strikes six at night, Doug wakes up and has the same kind of experience. Thus, you and Doug now have many more psychological connections than you had before, and you also aspire to build on these connections by compromising and coordinating much more often than you did before. But imagine that your psychological connectedness is still incomplete, in the sense that you still keep at least *some* beliefs, desires, and aims to yourself, and so does Doug.

Practically speaking, how does sharing a high degree of psychological connectedness affect your relationship with Doug? In part of the same way that sharing a chain of memory did: it makes it easier for you to live and let live. How? By ensuring that you have many of the same preferences about what to do. Before now you and Doug disagreed about pretty much everything: you wanted your shared life to revolve entirely around your poetry, and Doug wanted your shared life to revolve entirely around his music. But now that you share a high degree of psychological connectedness, you both feel differently. You really like the music that Doug makes, and you really like being able to help him make it. Similarly, Doug really likes the poetry that you write, and he really likes being able to help you write it. As a result, you both want the same kind of shared life now – a life where you get to write your poetry, Doug gets to make his music, and you each get to help the other do this, as well as experientially remember and anticipate his doing it.

In some respects, this changes your relationship with Doug significantly, since, insofar as you and Doug agree about what to do, you can rationally use first-personal methods *and*

strategies for interacting with each other. For example, suppose that, owing to your newfound psychological connectedness, you and Doug both want to go on a road trip, and you both agree on all the general parameters: where to go, when to go, how much money to spend, and so on. How should you proceed with respect to this issue? Easy. Just start planning your trip. For instance, you can form an intention to book a hotel tonight, with the expectation that Doug will remember your intention and act on it. Then Doug can book the hotel of his choice, because he knows that with respect to this issue, *his* decision is *your* decision. Then Doug can form an intention to book a rental car tomorrow, with the expectation that you will remember his intention and act on it. Then you can book the rental car of your choice, because you know that with respect to this issue, *your* decision is *his* decision. And so on. In short: when it comes to planning your trip, you and Doug can act as one. You can set ends for yourselves and then take the means to them, and you can do this without having to compromise at all. All you have to do is coordinate your behavior over time, by forming and acting on intentions.

In other respects, however, your relationship with Doug is exactly the same as it was before, since, insofar as you and Doug still disagree about what to do, you still have to use third-personal strategies for interacting with each other. For example, even though you and Doug agree that you should go on a trip, and even though you agree about the general parameters of the trip, you also disagree about many of the details. For example, Doug really wants to stay up late drinking the first night of the trip. But he knows that this will be impossible if you get an early start that morning, since anytime you get an early start, he has no energy left at the end of the day. Unfortunately, Doug also knows that, when morning rolls around, you will want to get an early start. So what should he do?

Well, Doug still has the same two options as always. First, he can try to compromise with you. For example, he can intend for you to sleep in tomorrow and then get an early start the next day, with the expectation that you will. Alternatively, if Doug wants to be able to stay out late drinking tomorrow without having to give up *anything* in exchange, then he can accomplish this by attacking and threatening you. For example, he can stay up all night drinking *tonight*, and then he can tell you, merely by framing the intention in his own mind, that if you even *think* of doing anything other than *exactly* what he tells you to, then he will personally see to it that you wake up with a hangover every day for the rest of your life.

What does morality have to say about this? Is Doug morally permitted to do this to you on the grounds that you now have the same preferences about what to do *with respect to other issues entirely*? Intuitively, the answer is still no. The fact that you now agree with Doug about what kind of hotel to book and what kind of car to rent *in no way* frees him to do whatever he wants to you otherwise. Granted, you are each now morally permitted to book the hotel of your choice and rent the car of your choice. But this is only because you have the same preferences with respect to these issues, and therefore you can do what you want without interfering with each other. However, if either of you were to change your mind about these issues, then you would once again have a duty to compromise. And you also, of course, have a duty to compromise with respect to the issues that you never agreed about in the first place. In short, then: you and Doug still have a duty to live and let live. Fortunately, insofar as you agree about how to live, you can do this without having to compromise as much as before. But insofar as you still disagree about how to live, you still have a duty to jointly commit to a fair set of laws, as well as a right to bind each other only insofar as you have to, in order to enforce these laws.

5. Sharing a self-narrative

Finally, imagine that, having spent a few weeks sharing a body and brain, a chain of memory, and a high degree of psychological connectedness, you and Doug start sharing a *self-narrative* as well. That is, whereas you previously used the first-person *plural* to refer to yourselves as a human being, you now use the first-person *singular* to do this. And whereas you previously used the first-person singular to refer to yourselves as *personalities*, you now use ‘my day self’ and ‘my night self’, or maybe even just ‘how I feel during the day’ and ‘how I feel during the night,’ respectively, to do this. Thus, when the clock strikes six in the morning, you wake up and not only *remember* everything that Doug thought, felt, and did over the past twelve hours, and not only *endorse* many of these thoughts, feelings, and actions, but also think of yourself as the *subject* of these thoughts and feelings and the *agent* of these actions. And when the clock strikes six at night, Doug wakes up and thinks about your actions and experiences in the same kind of way.

Practically speaking, how does sharing a self-narrative affect your relationship with Doug? In the same kind of way that sharing a chain of memory and a high degree of psychological connectedness did: it makes it easier for you and Doug to live and let live. How? By changing the *default unit of prudential concern* for both of you. Before this step in your fusion, you and Doug each identified with ‘me as a personality’ by default and with ‘us as a human being’ only with psychological effort. And you each thought that you rationally ought to maximize your preference satisfaction as a *personality* by default, and that you rationally ought to maximize your preference satisfaction as a *human being* if and only if you preferred to do this as a personality. But now that you share a self-narrative, all of this has changed. You each

identify with ‘me as a human being’ by default and with ‘my current personality’ only with psychological effort, and you each think that you rationally ought to maximize your preference satisfaction as a *human being* by default, and that you rationally ought to maximize your preference satisfaction as a *personality* if and only if you prefer to do this as a human being. And as a result, you and Doug will be more likely to agree about what to do now even when you have different beliefs, desires, preferences, etc. For example, when it comes to whether to drink tonight, you might both be inclined to think, “Well, *part of me* prefers not to drink at all and *part of me* prefers to drink a lot, but what *I* prefer, all things considered, is to drink a little.” And therefore you might both decide that drinking a little tonight is best. And insofar as you and Doug are both inclined to think and act this way, you will be able to rationally use first-person methods *and* strategies for interacting with each other *whether or not* you continue to have different beliefs, desires, preferences, etc. as personalities.

In other respects, however, your relationship is exactly the same as it was before. After all, even if you now identify with yourself as a human being by default, you can still identify with yourself as a personality if you want to. And insofar as you and Doug are still inclined to think and act in this latter way, you are still in the same exact situation as before, and you still have to use third-personal strategies for interacting with each other (even if you now think of these strategies in first-personal terms). For example, and using the language of your new self-narrative: Imagine that you are on a road trip, and you want to get an early start in the morning. But you know that this will be impossible if you drink tonight, and you also know that, when night rolls around, you will want to drink. What should you do?

Well, since you disagree with yourself across time about whether to drink (and since your presumption in favor of maximizing your preference satisfaction as a human being is, apparently,

not strong enough to resolve this disagreement), you still have the same two options as before. First, you can try to compromise with yourself. For example, you can intend to not drink at all tonight and then drink a lot tomorrow night, with the expectation that you will. Alternatively, if you want to get an early start tomorrow morning without having to give up *anything* in exchange, then you can attack and threaten yourself. For example, you can take all your beer and pour it down the drain, take all your money and give it to a friend to keep until morning, and then lock yourself in a room until your friend returns. Then you can tell yourself that if you so much as *think* about drinking tonight, you will do everything in your power to make sure that you stay sober the whole rest of the trip, and maybe for much longer than that.

What does morality have to say about this? To return to our previous way of talking: are you morally permitted to do this to Doug on the grounds that you now use different words to refer to each other, and you now identify with the human being as a whole by default? Intuitively, the answer is still no. The fact that you now use ‘me’ to refer to the human being as a whole and ‘my day self’ and ‘my night self’ to refer to you and Doug, respectively, does not free you to do whatever you want to each other. Similarly, the fact that you now identify with the human being as a whole by default does not free you to do whatever you want to each other when identifying with yourselves as personalities instead. In short: you and Doug still have a duty to live and let live. Fortunately, insofar as you want to maximize your preference satisfaction as a human being, you will be able to do this naturally, without having to think about it at all. But insofar as you feel inclined to maximize your preference satisfaction as a personality instead, you still have a duty to jointly commit to a fair set of laws, as well as a right to bind each other only insofar as you have to, in order to enforce those laws.

6. The circumstances of justice

In my view, this thought experiment supports, and finds support in, the view that, if x and y are in the circumstances of justice with each other, then they are different units of moral concern *whether or not* they are physically, psychologically, or narratively continuous with each other.

First, a bit of background. Rawls (1999) argues that the concept ‘justice’ applies to a particular group if and only if the *circumstances* of justice apply to that group. What are the circumstances of justice? They are the circumstances in which this virtue is possible to have, as well as necessary to try to have. So what are they in particular? Rawls provides a long list, which includes the following: the members of the group have to have different conceptions of the good, or in other words, different ends in life. They have to live in the same territory and use the same resources, often under conditions of scarcity. They have to have roughly similar mental and physical powers, and be motivated partly by self-interest and partly by benevolence. Finally, they have to be vulnerable to attack, from each other as well as from other groups. Rawls then claims that if these circumstances apply to a state, then the members of that state can, and must, find a way to live together in a just state. He also claims that the best way for them to do this is not by *eliminating* the circumstances of justice, but rather by *responding* to them in the right kind of way: by jointly committing to a fair set of laws, and binding each other only insofar as they have to in order to enforce these laws, so that they can live and let live in peace.^{vii}

This conception of the circumstances of justice explains and justifies our intuition that you and Doug still have duties of justice to each other at the end of fusion. After all, even though you have a much more intimate relationship now, you are still in the circumstances of justice with each other. Specifically, you still have different conceptions of the good. (You may agree

about how to live very often, but you still disagree about how to live *sometimes*, and this is enough for you to satisfy this condition.) You still live in the same territory (i.e., not just in the same nation, city, and room, but also in the same house, body, and mind) and use the same resources (i.e., not just resources like food, water, and electricity, but also resources like health, strength, and energy), often under conditions of scarcity. You are still roughly similar in physical and mental powers, and still motivated partly by self-interest (i.e., interest in promoting your welfare as a *personality*) and partly by benevolence (i.e., interest in promoting your welfare as a *human being*). Finally, you are still vulnerable to attack, both from each other and also, as a person, from other people. Thus, you can, and must, try to live together in a “just soul.” And as Rawls argues (and as I have argued here), the best way for you to do this is not by *eliminating* the circumstances of justice – for example by conditioning yourselves to always have the same preferences so that conflict between you never arises (after all, you might have good reason to want to stay psychologically fragmented to an extent) – but rather by *responding* to the circumstances of justice in the right kind of way: by jointly committing to a fair set of laws, and binding each other only insofar as you have to in order to enforce these laws, so that you can live and let live in peace.

Of course, there are a number of important differences between the just state and the just soul in practice, on this conception of the scope of morality. For example, given that most of us now live in nations with hundreds of millions of fellow citizens, practically speaking we have no choice but to compromise and coordinate with our fellow citizens *indirectly*, by selecting particular individuals to write, interpret, and enforce the laws that govern our interactions. In contrast, given that your day self and night self each have to make many decisions on behalf of the person as a whole each day, practically speaking they have no choice but to compromise and

coordinate *directly*, by thinking and acting in the spirit of compromise whenever possible. But as I have argued, none of this is morally relevant for present purposes. Our methods of interaction – our methods of writing, interpreting, and enforcing the laws that govern our interactions – might change as our group gets smaller and more intimate, but the goal is always the same: to find a set of laws that will help us to live and let live.

ⁱ As I will use these terms, x and y are physically continuous iff they share a body and brain, x and y are psychologically continuous iff they share a sufficiently high degree of psychological connectedness (i.e., a sufficiently high number of shared memories, beliefs, desires, aims, and so on), and x and y are narratively continuous iff they tell themselves a story according to which they are a single individual rather than separate individuals. See Eric Olson, *The Human Animal: Personal Identity Without Psychology* (Oxford: Oxford University Press, 1997) for a defense of a physical continuity theory of personal identity. See Derek Parfit, *Reasons and Persons* (Oxford: Oxford University Press, 1984) for a defense of a psychological continuity theory of personal identity and what matters in morality. And see Marya Schechtman, *The Constitution of Selves* (Ithaca: Cornell University Press, 1996) for a defense of a narrative continuity theory of personal identity and what matters in morality.

ⁱⁱ Parfit 1984, pp. 321-50. To be clear, I will not be taking a stand in this paper on whether the relation ‘what matters in morality’ has the logic of numerical identity.

ⁱⁱⁱ For more on the circumstances of justice, see David Hume, *An Enquiry Concerning the Principles of Morals* Section III Part 1 (Oxford: Oxford University Press, 1998) and John Rawls, *A Theory of Justice: Revised Edition* (Cambridge: Harvard University Press, 1999).

^{iv} For an overview of the debate about whether persons endure or perdure, see Katherine Hawley, "Temporal Parts", *The Stanford Encyclopedia of Philosophy* (Winter 2010 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/win2010/entries/temporal-parts/>>. And for an overview of the debate about whether persons can have multiple selves, see Jennifer Radden, "Multiple selves," in *The Oxford Handbook of the Self*, S. Gallagher (ed.) (Oxford: Oxford University Press, 2011).

^v Throughout this thought experiment, I will use ‘you’ to refer to the personality that controls this new body from 6am to 6pm and ‘Doug’ to refer to the personality that controls this new body from 6pm to 6am. I will be using these names only for ease of exposition; nothing in my argument will rest on the metaphysical view that you and Doug survive this fusion.

^{vi} Some philosophers use ‘quasi-memory’ to refer to a state which is neutral with respect to whether the remembering subject is the same person as the subject whose experience is being remembered. So, if you prefer, you can read me as saying that you and Doug share quasi-memories instead of memories. For more on quasi-memory, see Sydney Shoemaker, "Persons and Their Pasts," *American Philosophical Quarterly* 7, 1970.

^{vii} Rawls (1999), pp. 109-112.