



# Identity and influence

Daniela Dover<sup>1</sup>

Received: 28 March 2023 / Accepted: 20 July 2023 / Published online: 15 November 2023  
© The Author(s) 2023

## Abstract

How worried should we be about how impressionable we are—how susceptible we are to being influenced and even transformed by our encounters with one another? Some moral philosophers think we should be quite worried indeed: they hold that interpersonal influence is an especially morally dangerous way to change. It calls for additional moral scrutiny as compared with vectors of change that come from within the influencee’s own psyche—their antecedent values, desires, commitments, and so forth—just because it has an external source. I argue that this heightened scrutiny of exogenous sources of change is unwarranted. Dramatic psychic changes do call for reflection and critical scrutiny, especially when they are sudden. But this scrutiny need not be concerned with the procedural issue of whether the impetus for the change came from inside or outside the changing person’s antecedent psychology. We can just evaluate the substantive changes themselves.

**Keywords** Influence · Identity · Authenticity · Self · Transformative experience

How worried should we be about how impressionable we are—how susceptible we are to being influenced and even transformed by our encounters with one another? Some moral philosophers think we should be quite worried indeed: they hold that interpersonal influence is an especially morally dangerous way to change. It calls for additional moral scrutiny as compared with vectors of change that come from within the influencee’s own psyche—their antecedent values, desires, commitments, and so forth—just because it has an external source.<sup>1</sup> Call this the *heightened scrutiny*

---

<sup>1</sup> You may already be asking yourself how we can tell ‘external’ from ‘internal’ sources of psychic transformation. I will return to this question later. For now, I will grant for the sake of argument that some intuitive distinction in this neighborhood can hold up and focus on whether that distinction matters morally. Heightened scrutinizers think it does, and I do not.

---

T.C.: Transformative Experience, Authenticity, and Rationality

✉ Daniela Dover  
daniela.dover@gmail.com

<sup>1</sup> Oxford University, Sidney House, Flat 60, Old Ford Road, London E29QB, England, UK

*view*. I want to reject the heightened scrutiny view. Dramatic psychic changes do call for reflection and critical scrutiny, especially when they are sudden. But this scrutiny need not be concerned with the procedural issue of whether the impetus for the change came from inside or outside the changing person's antecedent psychology. We can just evaluate the substantive changes themselves.

## 1 Problems of transformation, new and old

Consider the phenomenon that Laurie Paul calls 'transformative experience': experiences that 'teach us things we cannot know about from any other source but the experience itself' (p. 3) and through which we can expect to become 'an unknown, dramatically changed, new self' (p. 51).<sup>2</sup> Paul tends to focus on rites of passage that people can (ostensibly) 'choose between having [and] avoiding' (p. 122), such as embarking on a career, getting married, or becoming a parent. Cases of interpersonal influence—falling in love, working with a great teacher, falling under the sway of a charismatic leader, or finding a new philosophical interlocutor—are less likely to show up in a short obituary. But they can change us just as profoundly.

Paul is concerned with how we might approach choices about whether to undergo a transformative experience. She worries that transformative choices pose problems for standard models of rational decision-making, for two reasons. First, these are experiences whose subjective character you cannot anticipate until you get there. You can hear other people talk about how much they love their children, but you will not really know what having children will be like for you until you have done it. This makes it hard to set about maximizing expected value by picking the course of action that is most likely to satisfy your preferences.

Moreover, even if you could anticipate how likely having a child would be to satisfy your current preferences, it would remain unclear whether that is even what should matter here, since having a child is just the sort of thing that is likely to change your preferences. Perhaps now you prefer to stay out late drinking and sleep until noon; having a child would put a damper on that. But for all you know, once you have the child, you will love nothing more than getting up early and taking them out for a wholesome morning walk in the park. So whose preferences matter—those of your current boozy night-owl self, or those of the sober lark self that might emerge after your child does?

So far, these sound like prudential problems. We cannot use some of the rational tools we might normally use to get what we want. And yet Paul ends up framing her concerns not only in terms of the normativity of prudential rationality but also in terms of the unabashedly ethical concept of *authenticity*. However startling this appeal to authenticity might be for a reader expecting an intervention in decision theory, Paul is absolutely right to explore the ethical roots of her concerns about transformation. It is prudential rationality that tells me how to go about trying to get what I want. But prudential rationality is silent as to why it should be so important to work so hard—perhaps even breaking out a calculator—to figure out what I want in the first

<sup>2</sup> All references to Paul are to Paul (2014).

place and try to get it (let alone to maximize how much of it I get). Why should that be what governs my life? As anyone who has taken an ethics class has been reminded, I could instead try to figure out how to maximize the good, or to discern and carry out my duties. Less nobly, I could look at what everybody else seems to want for themselves, and try to get some of what they are having. Or I could go by what my mother wants for me. Or I could simply do whatever seems easiest at the moment. So there is an irreducibly ethical idea lurking behind Paul's assumption that it is our own preferences—or values, to put it more respectably—that we should try to discern and realize.

Here is some of the language Paul uses to invoke this idea:

You are expected to take charge of your own destiny. You chart your future, deliberating and reflecting on who you really are and what you really want from life.... You live an authentic life by faithfully modeling your preferences, and you live a rational life by matching your choices to these preferences. Rational authenticity, then, is hewing as close as you can to the kind of life that best realizes your dreams, hopes, and aspirations. (p. 105)

We need not worry too much about the details of this particular formulation. Indeed, 'authenticity' is just one among many possible concepts Paul could have reached for to call these ideas to mind. The language of autonomy, sincerity, integrity, steadfastness, wholeheartedness, self-authorship, self-ownership, self-discovery, self-reliance, self-mastery, self-determination, or self-realization can conjure up more or less the same set of notions.<sup>3</sup> Moral theorists have developed rich conceptions of every one of these concepts, but it is the cluster of vague notions that I want to focus on here: without getting too hung up on terminology, we can all recognize the general outlines of the moral vision Paul is invoking. According to this image of ethical life, what it is to be good at living is to be good at taking charge of your own destiny and pursuing your own dreams.

We can think of this cluster of ideas as an *ethos* of authenticity. I call it an 'ethos' (rather than, e.g., a 'norm' or an 'ideal') for several reasons. First, I want to draw our attention to the fact that the ethos of authenticity is not an action-guiding norm, nor even a well-specified end or aim. It is so far no more than a collection of hazy representations whose precise normative implications, if any, remain to be specified. Second, unlike 'ideal', the term 'ethos' (like, e.g., 'sensitivity') remains explicitly neutral as to the normative status of the phenomenon to which it refers. Third, calling it an 'ethos' rather than an 'ideal' draws our attention more explicitly to the fact that such collections of representations include both positive and negative elements. In other words, an 'ethos' includes both representations of ideals, bathed in a halo of cultural approbation, and their corresponding pejorative representations, shrouded in

<sup>3</sup> Autonomy obviously stands out in this list for being the ethical concept that has been most systematically theorized in modern Western philosophy. One of the theorists I discuss below, Thomas Hill, is a card-carrying neo-Kantian. But many moral philosophers who invoke this ethos, like Paul (and Paul's discussion is rare in acknowledging this), are simply browsing in a curiosity shop of inherited values (the 'culture's ordinary way of thinking', as Paul [85] puts it). Short of developing or accepting a moral theory as systematic as Kant's, perhaps this is the best any of us can do. All the more reason to look closely and critically at the wares.

cultural opprobrium. The ethos of authenticity, as we shall see, functions not only to lionize certain ways of feeling, thinking, and living but also to stigmatize others.

It is difficult to subject an ‘ethos’ to philosophical criticism. When defending positions that align comfortably with an ethos that is taken for granted by their home culture, philosophers are less likely to argue for, or even clearly to state, the presuppositions that a reader skeptical of the ethos would want to call into question. Moreover, a culturally dominant ethos, such as the ethos of authenticity, straddles philosophical and everyday moral discourse in ways that differ from how a more antinomian ethos, such as the anarchist ethos of mutual aid, or the Quaker ethos of nonviolence, would. As Richard Moran (2011, p. 182) notes in a discussion of Iris Murdoch’s engagement with existentialism, Murdoch sought to reject ‘a conception of ourselves which can be found in movies, popular songs, and forms of romantic life as much as in works of philosophy’, and which ‘exists as a style and an attitude much more than simply a theory’. The challenges of such a task become even more acute when one is working within the genre conventions of analytic philosophy, which call for sharply-defined targets of critique.<sup>4</sup> Despite these challenges, the task is nonetheless worthwhile—arguably much more so than the task of criticizing sharply-defined views held only by professional philosophers.

Paul’s discussion endorses but does not defend the ethos of authenticity. Instead, with refreshing frankness, rather than treating the ethos as uncontroversially appealing, she simply notes how compelling it is likely to be for decision-makers who are ‘member[s] of an affluent, contemporary Western culture’ (Paul, 2014, p. 85)—which is to say, for almost all of her likely readers. What Paul is saying is that, *if* you are committed to such an ethos, *then* transformative choices are especially problematic: it is hard to know how to decide in a way that is both rational and authentic.

Now I want to call attention to another, more subtle but equally important idea lurking in the background of Paul’s discussion. This is a certain implicit picture of the human self (or of ‘identity’—I will use these terms interchangeably here). Paul conceives of the self as constituted by a person’s core preferences and point of view (p. 16), but that is not the assumption that I am talking about. I am talking about a more general assumption that runs upstream of this preference-based theory of the self, which is the idea that we are transformable creatures in the first place.

This is not just the idea that we can change. Everything changes, especially living things. But there’s change and then there’s change. Nobody talks about ‘transformation’ in the context of trimming your nails or improving your golf swing. To talk about ‘transformation’ is to draw a distinction between two kinds of change: a mundane kind that does not merit the term ‘transformation’ and a more dramatic one that does. To call an experience like having a child ‘transformative’ says two things about us. First, human beings are the sort of thing that can undergo the latter, more earthshaking sort of change. Second, this does not require any dramatic scenario of the sort to be found in the personal identity literature after Parfit (1984). I need not undergo science-fictional fission or fall into the hands of a malicious brain surgeon. I am the sort of creature that can be ‘transformed’ through experiences that, while they will stand out as plot

---

<sup>4</sup> For an example of how to criticize an ethos while working outside those genre conventions, look at Nietzsche.

points in the context of an individual human life, are utterly commonplace in the grand scheme of things—like love, grief, violence, illness, gain or loss of faith.

This basic notion—that we are so transformable that the ordinary vicissitudes of human life can change our ‘selves’ or ‘identities’ (whatever those are) so profoundly as to tempt us and others to say that we are ‘no longer the same person’—is more unsettling the more deeply one is committed to the ethos of authenticity. That is because, in its crudest forms, what that ethos is saying to us is: Stay the same. Do not change. Or, in a more sophisticated variation: change, but only in ways that make you into a fuller realization of what you already are—as an anonymous referee for *Synthese* put it to me, ‘change in such a way that you become the “truer” version of yourself’.<sup>5</sup> Which is to say: don’t *really* change, just self-actualize.

Now, of course, no philosopher is going to come out and officially defend the claim that, while people do tend to change, they should generally try their best not to. That would be a strawmanly position to attribute to most philosophers who are drawn to the ethos of authenticity—though some do come surprisingly close.<sup>6</sup> And yet consider some of the everyday idioms we reach for when we want to invoke this ethos, idioms of the sort to which philosophers invariably turn, sooner or later, in theorizing it. ‘Stay true to yourself’, ‘don’t lose yourself’: such clichés counsel us to *preserve* something, be it our identities, our plans, our values, or our commitments, across time. From such a preservationist perspective, the way that life has of transforming us can be profoundly disturbing. How can we figure out who we really are, chart our destinies, and realize our dreams, given how vulnerable we are to being changed just by the ordinary things that life throws at us?

Paul is not the first to have this worry. There is a very long line of philosophers who have articulated a concern about how our malleability can undermine our authenticity.<sup>7</sup> Whether they use the term ‘authenticity’ or whether instead they appeal to one or more of the kindred notions listed above, Western philosophers have been particularly concerned, for the past three and a half centuries or so, not only with how we can remain just or virtuous (those were central concerns already for Plato and Aristotle) but also, and often even primarily, with something much more reflexive than that: How, in the face of inner and outer changes, can we remain true to ourselves?

This question has been taken up explicitly by moral philosophers at least since Rousseau. It was also a central concern for the author(s) of the inspirational quotations printed on the tags dangling out of the many cups of tea I drank while working on this paper, which often seemed to be trolling me with their enthusiasm for the ethos of authenticity. Some offered admirably succinct summaries of the advice I was

<sup>5</sup> Thanks to Samuel Reis-Dennis for discussion of this point.

<sup>6</sup> Here is Jennifer Hawkins (2019, p. 252), for instance, just after acknowledging the ‘well-known fact that people change over time’: ‘I take it that most really radical change is bad for the person who undergoes it. Consider Phineas Gage, the nineteenth-century railroad worker...’ (It is telling just how often Phineas Gage’s traumatic brain injury—his brain was pierced straight through by an iron rod—turns up in this literature as an example of psychological change.)

<sup>7</sup> There have been many excellent intellectual-historical accounts of this tendency; see especially Berman ([1970] 2009), Starobinski (1988), and Trilling (1972). Charles Taylor (1992) covers some of the same ground in the context of a survey of the ‘malaises of modernity’; while I cannot share the nostalgia for premodern ethical traditions that animates Taylor’s discussion, I am in sympathy with many of his criticisms of what he calls the ‘ethics of authenticity’.

encountering in more august venues: ‘master the unknown by knowing your deep self’ (YogiTea, 2023a).<sup>8</sup> Paul’s work on transformative experience interests me as a novel formulation of this familiar, centuries-old ethical preoccupation.

## 2 Influence as thralldom: motivating heightened scrutiny

Authenticity-based concerns can arise in the context of all sorts of change. People might worry about converting to a new religion, or becoming obsessed with a new self-help book, or taking psychedelic drugs, or psychiatric ones: these are all contexts where you, or the people who love you, might worry that you could ‘lose yourself’. But interpersonal influence is a vector or mechanism of change that has often been thought to raise especially acute authenticity-based concerns. This section outlines some of these concerns and begins to indicate why I find them unpersuasive.

Influence is ubiquitous in human life. As children and adolescents, we are invariably influenced by various caretakers, teachers, and peers. But I restrict my focus here to the sorts of influence that can happen even well into adulthood, such as when we fall in love, change our minds after an argument, or admire someone so much that it gives us appreciation for a new sort of art, hope for a revolution, or faith in a new god or a new philosophical theory. Taking our cue from Paul, call these cases of *transformative interpersonal influence* (‘influence’ for short).<sup>9</sup> Philosophers who favor heightened scrutiny are especially worried about how we can preserve our authenticity in the face of such influence. They counsel us to conduct background checks on the sources of our transformations and regard (what we take to be) the exogenous ones with extra suspicion.

Take, for instance, a stimulating recent paper by Farbod Akhlaghi (2022) that approaches the problem from the perspective of the would-be influencer. Akhlaghi argues that when our friends are facing a transformative choice, unless they proactively come to us for advice, we have a standing duty to refrain from influencing their decisions. This is because our friends have a ‘moral right’ to make these choices on their own, a duty which Akhlaghi derives from ‘the moral value of autonomous self-making’ (2022, p. 9) or ‘self-authorship’ (6 ff.). Making or authoring ourselves, he thinks, is ‘crucial for us and others to see ourselves as *ourselves*’ (6, emphasis in original). Although Akhlaghi never uses the word ‘authenticity’, he is drawing on a similar set of cultural tropes and values as Paul. But what Paul frames as a problem faced by solitary rational agents who want to preserve their authenticity while charting their own courses through uncertainty and upheaval is framed by Akhlaghi as problem for moral agents who want to respect their loved ones’ solitude as they do just that.

Views like Akhlaghi’s owe their initial plausibility to their close connection with the ethos of authenticity. Authenticity, as we have seen, is often praised as self-authorship, self-ownership, self-discovery, self-mastery, and self-realization. Intuitively, you will

<sup>8</sup> But cf. YogiTea 2023b, which counsels that ‘the unknown has to be trusted’. In what follows, all quotations from teabag tags are from batches of ‘YogiTea’ purchased in London in February 2023.

<sup>9</sup> From here on out, the word ‘influence’ will always refer to cases in which someone is influenced, in adulthood, by a single other person or small group of other people known to them personally—not by large group phenomena or by a generic ‘das Man’.

not want others messing around with all these reflexive processes of self-selfing. Instead, ‘you should develop mental strength’ (YogiTea, 2023c) by ‘get[ting] in touch with yourself’ (2023d), since ‘your deep inner self is your brightest light’ (2023e).<sup>10</sup> But the most influential discussions favoring heightened scrutiny also use memorable intuition-pumping cases to motivate it. Unsurprisingly, these discussions often focus on especially malign-sounding cases of influence. Akhlaghi’s three central examples are trying to get your friend *not* to quit his lucrative job to become a schoolteacher, trying get your brother *not* to become a parent, and trying to get your high-school sweetheart *not* to go to college. (The last case adds a gendered dimension, with ‘Jack’ trying to convince ‘Jill’ to turn down a full scholarship.) Trusting his readers to root for public service, natality, and higher education, Akhlaghi puts these examples forward in the service of an argument that influence as such is morally suspect on *procedural* grounds.<sup>11</sup>

Thomas Hill pursues a similar strategy in an enormously influential paper that has given rise to an extensive literature in the fifty years since its publication. Hill’s case of the ‘Deferential Wife’ is meant to illustrate a ‘submissive attitude’ (Hill, 1973, p. 90) that he calls ‘servility’, and which, he argues, ‘betrays the absence of a certain kind of self-respect’ (p. 89). A patriarch’s fantasy bride, the Deferential Wife understands and values herself, and expects to be understood and valued by her husband, exclusively in terms of ready-made social scripts of wifely virtue. She is so ‘utterly devoted to serving her husband’ that she ‘buys the clothes he prefers, invites the guests he wants to entertain, and makes love whenever he is in the mood’ (p. 89). She happily moves for his job, since she counts ‘her own friendships and geographical preferences insignificant by comparison’ to his (p. 89). She does not only ‘defer to her husband in certain spheres’ or ‘as a trade-off for his deference’ in others (p. 89). ‘On the contrary, she tends not to form her own interests, values, and ideals [at all] and when she does, she counts them as less important than her husband’s’ (p. 89).

Several features of this description should give us pause. Written in 1973, at the height of the women’s liberation movement in the US, it was doubtless intended as a sympathetic portrayal of a life warped by internalized oppression. But Hill’s interest in servility as a ‘vice’ or ‘moral defect’ injects a note of scorn that gets in the way

<sup>10</sup> According to YogiTea, these are necessary conditions not only for authenticity but also for felicity: ‘without realising who you are, happiness cannot come to you’ (2023f).

<sup>11</sup> Akhlaghi also uses especially ominous-sounding abstract nouns and verbs to refer to the phenomena of influence. The neutral term ‘influence’ appears only once in the paper, and the only slightly more nervous-making ‘advice/advise’ and ‘convince’ are entirely absent. Instead, we have ‘stop’ (four times) and ‘interfere/interference’ (38 times). Akhlaghi explicitly includes under the rubric of ‘interference’ the strategy of rationally persuading our friends by appeal to what we know about their values—a practice one would not reach for the words ‘stop’ or ‘interfere’ to describe unless one were already committed to heightened scrutiny. Such rhetorical gambits are ubiquitous—and often, I think, entirely legitimate—in philosophical writing. I pick on Akhlaghi here not because he relies more heavily on them than others do but because his rhetoric rows downstream: it is being used to uphold familiar moral tenets that readers could easily have been recruited to reaffirm by means of more measured language. We tend to notice rhetoric most when we find the claims being made least persuasive. By the same token, it is easy to miss even the most potent rhetoric when it reinforces moral common sense. Insofar as moral philosophers are concerned to scrutinize contemporary moral norms, it is especially important to attend closely to metaphor and diction when they are employed in defense of the cultural status quo.



of the realism of the sketch.<sup>12</sup> When Hill says that the wife ‘tends not to form her own interests, values, and ideals’, we should stop to wonder what interests, values, and ideals, if any, this woman had before she met her husband. We are told she has no real friends. But surely she was close to someone at some point in her maiden days. Does she admire her father? (If not, then how did she get into the habit of worshiping authoritative men so uncritically?) Does she admire her mother? (If not, it is hard to imagine where she learned to be such a paragon of wifely virtue.) Speaking of parents, does she care about her own children? (Surely she has children.) If so, does she really never disagree with her husband about how to raise them? And if she has scarcely any ‘interests, values, and ideals of her own’, then how does she even manage to have ‘friendships and geographical preferences’ that could threaten to interfere with her husband’s career prospects? Real lives, even under circumstances of severe oppression, are more variegated than this.<sup>13</sup>

Even supposing that Hill’s description of the Deferential Wife could be a fair summary of an actual human moral psychology, it would be hard to discern precisely which of that person’s problems resulted from her susceptibility to outside influence. There is really quite a lot going wrong in this woman’s life, substantively speaking. For instance, we have already heard enough to guess that she spends a lot of her time wearing uncomfortable clothes (‘the clothes he prefers’) while entertaining boring guests (‘the guests he wants to entertain’) in a new city that does not feel like home. It hardly matters whether these boring guests include her own friends, since she has only the most tenuous and disposable of friendships. These dreary evenings are often capped off by ‘making love’ when she is not ‘in the mood’—which is to say, by very bad sex. Even if, as Hill rather implausibly stipulates, this woman is nonetheless ‘happy’ and ‘loves her husband’, these are straightforward respects in which her life is pretty lousy, substantively speaking. They make it hard to say how much of the pity and revulsion she is meant to invoke in the reader is due to her susceptibility to influence (a feature of *how* her identity has been shaped) and how much is due to *what* is shaping it (the substantive character of the influences she has absorbed). In short, caricatures like Akhlaghi’s and Hill’s strike me as misleading when used, as they so often are, as entrees into a philosophical discussion of the ethics of influence.

I have focused on Akhlaghi and Hill so far as especially clear cases of the phenomenon that interests me—namely, philosophers relying (more or less explicitly) on the ethos of authenticity in defending heightened scrutiny. But Akhlaghi and Hill are not alone here. A closely related proposal is offered, for instance, by Michael Garnett (2023) who draws on Hill’s Deferential Wife case to argue that the wife’s ‘willingness

<sup>12</sup> For the claim that servility is ‘as much a vice as arrogance’, see Hill (1973, p. 88). For servility as a moral ‘defect’ see Hill (1973, pp. 87, 92, 93, 94, 97, and 100).

<sup>13</sup> In the text I have emphasized that the story fails to persuade because it implausibly lacks conflict. On the other hand, its plausibility *as a story of oppression* depends on the possibility of the very conflict it suppresses. My complaint here is not about the use of toy examples like this; I am saying that the example fails by the lights of that genre. The perfect analytic philosophy toy example contains all and only the details that are relevant to the question. My point is that this one leaves out details that would have to be included in a fair sketch of any real (or believable fictional) person designed to bring out the features most relevant to this question. For a different, and very powerful, critique of Hill’s use of examples, see Boxill and Boxill 2015.



to do whatever her husband wants just because he wants it is key to, indeed constitutive of, her deference' (p. 198). Like Hill's, Garnett's diagnosis of the Wife focuses on the source, not the content, of her motivations: her problem is that 'his preferences simply are her reasons' (p. 198). Guarding against deference of this kind will require heightened scrutiny of (ostensibly) reason-giving desires when, and just because, they have an exogenous source.

We can see more clearly why Garnett's conception of deference recommends heightened scrutiny when we compare his view with that of a neo-Kantian like Christine Korsgaard. Korsgaard (1999, pp. 16–17) discusses a case, adapted from Jane Austen's *Emma*, which is very similar to that of the Deferential Wife, except that the background power differential—a political detail that is glossed over in Korsgaard's discussion, as it is in Hill's—is one of class rather than gender. The pliable Harriet makes all her major decisions by asking Emma what she should do, and then doing just that. According to Korsgaard, Harriet's action is 'defective as autonomous action' (p. 16, emphasis in original), and Harriet herself 'is heteronymous', because she 'allows herself to be governed in her choices by a law outside of herself' (p. 16). For Korsgaard, however, crucially, for an action to come outside of *oneself* means for it to come outside of one's *rational will*, not outside of one's psychology taken as a whole. Were Harriet were to allow herself to be governed in her choices by her own appetites, this would constitute heteronomy in the very same way, and for the same reasons, as allowing herself to be governed by Emma's whims. For, according to Korsgaard's moral psychology, Harriet's own inclinations are just as 'external' to Harriet's rational will as Emma's inclinations are: when I akratically smoke a cigarette just because I feel like it, I am governed by a law 'outside of myself' just as much, and in the same way, as when I deferentially smoke a cigarette just because you offered me one.

By contrast, when, in defining the notion of deference, Garnett writes that 'B makes a deferential choice...just in case B takes *some A's* wanting it to be the case...to be an especially weighty reason for choosing' (p. 208, emphasis added), it matters to him that A is not the same person as B. What makes B's choice deferential is that its source lies in another person. To avoid such deference—and ensure that our lives will be 'self-authored' rather than 'ghostwritten' (p. 213)—we would have to adopt a policy of being more skeptical of exogenous sources of motivation than we are of sources that were already present within us before our encounter with this particular 'other'. In short, Garnett's view implicitly recommends a commitment to heightened scrutiny as a prophylactic against deference.

Another related proposal can be found in Bernard Williams' canonical work on internal and external reasons (1981, 1995, 2001, 2006a), which relies upon—and, in a series of papers written over two decades, theoretically elaborates—a contrast between

exogenous and endogenous sources of influence.<sup>14</sup> In ‘Values, Reasons, and the Theory of Persuasion’ (2006a, 115ff.), Williams discusses the ‘constraints’ that must be observed by a speaker who wishes to afford ‘deliberative assistance’—what Williams elsewhere (1995, p. 36) calls ‘advice in the “if I were you...” mode’—rather than effecting a transformation that comes ‘closer to conversion’ (2006a, p. 115). Such a speaker must confine herself to citing considerations whose force the addressee could be led to acknowledge via a ‘sound deliberative route’ from their own antecedent desires, values and commitments (their ‘S’, short for ‘subjective motivational set’).<sup>15</sup> Although Williams largely keeps his discussions of this contrast at a metaethical level, avoiding pronouncements about the morality of transformative ‘conversion’ as opposed to ‘sympathetic advi[ce]’ (2006a, p. 117), numerous moments in his work—for instance, his talk of the risk of ‘coerci[on]’ (2006a, p. 118) involved in the former—strongly suggest that, like Akhlaghi, Garnett, and Hill, Williams saw this contrast as ethically, not just metaethically, significant. Indeed, though he remained cagey about this, it is, I think, fair to attribute to Williams the view that a speaker who wishes to treat their addressee with consideration and respect *should*, at least in most cases, stick to helping the addressee to think through the implications of their antecedent *S*.<sup>16</sup>

George Tsai (2014) goes further. For Tsai, even a speaker who scrupulously restricts herself to appealing to considerations that were already important to her interlocutor still risks interfering with the interlocutor’s right ‘to be the one who makes her own life, to be the author of it’ (Tsai, 2014, p. 79).<sup>17</sup> Say that A tries to persuade B to

<sup>14</sup> Williams was always careful to acknowledge how difficult this distinction is to draw in a principled way. As early as ‘Internal and External Reasons’, he emphasized that ‘there are no fixed boundaries on the continuum from rational thought to inspiration and conversion’ (1981, p. 110). Williams’ understanding of the contrast between helping one’s interlocutor to think through their own reasons, on the one hand, and bombarding them with external [putative] reasons, on the other, clearly evolved in various ways in the two decades following the original 1980 publication of ‘Internal and External Reasons’. I do not have space here to go into exegetical questions about the course of this evolution. What is important to me in the text that follows is that, throughout this period, Williams thought the contrast exceedingly important and expressed—or, at the very least, evinced—some negative normative attitudes toward what he saw as hubristic or hectoring attempts to bring about, in another person, transformations lying closer to the ‘conversion’ end of the spectrum. See n. 16 below.

<sup>15</sup> An agent’s *S* can include ‘desires’ as well as ‘dispositions of evaluation, patterns of emotional reaction, personal loyalties, and various projects, as they may be abstractly called, embodying commitments of the agent’ (Williams 1981, p. 105). For the evolution of Williams’ notion of a ‘sound deliberative route’, see his 1995, p. 35ff.; 2001, p. 99ff.; 2006a, p. 110 ff. (But note that ‘Values, Reasons, and the Theory of Persuasion’ (Williams, 2006a) was originally published in 1995.)

<sup>16</sup> Williams suggests that appealing to internal reasons is the cooperative thing to do when he calls the interlocutor who sticks to citing internal reasons an ‘advisor or fellow deliberator’ (2006a, p. 115) who is ‘helpful’ (p. 115) and ‘scrupulous[]’ (p. 116) as well as ‘sympathetic’ (p. 117). In his last word on the subject (see n. 15 on the order of composition), he warns that externalist claims risk amounting to no more than ‘bluff and brow-beating’ (2001, p. 95). And of course the term ‘conversion’ itself, in such a resolutely secular thinker—one who is, moreover, famously critical of the influence of Christianity on contemporary morality (2006b)—has a certain whiff.

<sup>17</sup> Tsai clearly believes that even citing what Williams would call ‘internal reasons’—indeed, even citing what one’s interlocutor *already recognizes* as internal reasons (as opposed to what they might be brought, via a sound deliberative route, to recognize as such)—runs a risk of objectionable paternalism. This comes through, for instance, when Tsai clarifies that we ‘can resent [persuasive] interventions...*even when we perfectly well see*’ that the content of the ‘considerations offered to us has the status of good reasons’ (2014, p. 91, emphasis added).

eat a bowl of pasta by noting that B is currently trying very hard to summit Mount Whitney, that B is prone to altitude sickness, and that eating extra carbohydrates helps to ward off altitude sickness. Clearly, rather than ‘merely browbeating’ B, pummeling her with considerations that have no foothold in her subjective motivational set, A is simply helping B to think through how to get what B already, antecedently and independently of A, badly wants. For Tsai, even such a light touch might, in theory, bring it about that B’s subsequent decision ‘cannot...be seen as flowing from the *independent exercise* of her capacities to canvass and weigh reasons’ (p. 79, emphasis added). This in turn risks subjecting B to ‘a form of loss in being unable to see herself fully as a self-directing agent realizing purposes of her own’ (p. 79). For Tsai, self-authorship requires heightened scrutiny not only of novel incursions into a person’s S, but also of the involvement of outsiders even in the processes by which a person ‘canvass[es] and weigh[s]’ their own, home-grown psychic cargo.

Akhlaghi, Garnett, Hill, Tsai, and Williams have all arrived, by different routes, at more or less extreme versions of the suggestion that we should be especially careful to guard against outside influence as a source of change in the self—more careful than we are to critically scrutinize changes that arise from within. I am glossing over many important and interesting differences among their views, each of which would repay far more sustained attention than I can give it here. What is salient for my purposes is that all of these authors are concerned to impose especially stringent requirements on interpersonal influence—as opposed to ‘internal’ contents, forces, and processes—as sources of psychic change and motion. My strategy from here on out will be to offer the reader an alternative way of thinking about how the influence of others shapes us, one that does not stack the deck in favor of subjecting these influences to especially rigorous scrutiny.

### 3 How to have friends

We are all born into a web of social relationships, and these early formative relationships are all unchosen. You are reading this today because at least one person took care of you for quite a few years; this was a person whom you did not choose, and who did not choose you (even if they chose to conceive a child, they did not pick you out of a lineup of possible children). These early non-elective caring relationships are often far from idyllic. But even the best of them typically cannot be all that psychically sustains us for the rest of our lives. For most people, a full adult life has to include at least a few *elective* relationships with people who have picked us out of the crowd, choosing to devote their finite time, care, and attention to us in preference to other Others whom they might—entirely blamelessly—have preferred. Most of us cannot get by without friends, or lovers, or housemates, or bandmates, or all of the above (from here on I will ignore differences among genres of elective relationships and speak of them all as ‘friendships’).<sup>18</sup>

<sup>18</sup> As I note elsewhere (Dover 2022, p. 198), friendship’s status as an elective relationship with low ‘extra-relational stakes’ as compared to, e.g., marriage—what we lose when a friendship ends is often ‘just’ the friendship itself—makes it an especially revealing case for the philosophical study of interpersonal dynamics.

The ethical significance of this familiar fact—that we cannot get by without elective relationships—is often overlooked in philosophical discussions of interpersonal relationships that focus on explicitly moral interpersonal attitudes such as respect, recognition, concern, and benevolence, along with the reactive attitudes. In his landmark paper ‘Freedom and Resentment’, P.F. Strawson emphasized ‘how much we actually mind, how much it matters to us, whether the actions of other people—and particularly of some other people—reflect attitudes towards us of goodwill, affection, or esteem on the one hand or contempt, indifference, or malevolence on the other’ (Strawson, 2008, pp. 5–6). But he went on to devote much more attention to the moralized attitudes on this list—goodwill and malevolence—than he did to affection, esteem, indifference, or contempt. Many subsequent discussions of interpersonal relationships have done the same, stressing our needs and entitlements to be respected and recognized while not at the same time acknowledging the fact that we also want and need attitudes from others to which we are *not* entitled—such as to be, for instance, liked, or loved.<sup>19</sup> And yet these more discretionary attitudes—the ones, like affection, that we cannot, in adulthood, demand from one another as a matter of right—are just as necessary to elective relationships as are more morally regulated attitudes like benevolence and ‘recognition respect’ (Darwall 1977). In this section I consider two central challenges that we face in our attempts to establish and maintain elective relationships and reflect on some of the ways in which these challenges shape and transform us.

First, we have to become and remain sufficiently *intelligible* to each of our various ‘others’, respectively. As David Velleman (2009) argues, individual agents function best when they have a workable—which need not be to say entirely veridical—conception of themselves and of what they are up to. What holds for individual agency here holds for relationality as well. Just as we need to be sufficiently intelligible to ourselves to function as agents, we need to be sufficiently intelligible to others to function as partners in ongoing relationships.<sup>20</sup> When others cannot make sufficient sense of us—in an ordinary folk-psychological sense of ‘sense-making’—it becomes difficult for them to relate to us. So just as agents have reason to strive for first-personal intelligibility in order to function well in their individual endeavors, they have reason to strive for second-personal intelligibility—not intelligibility as such (if there is any such thing), but intelligibility to a *particular* other person—in order to function well in a relationship with that person. This means that insofar as we want to be able to have an ongoing relationship with another person, we have reason to cooperate in that person’s attempts to make sense of us.

But being intelligible to another person is not yet enough to establish or sustain an elective relationship with them. An elective relationship functions only if both partners want to engage in it. This presents us with a second challenge: we have to become and

<sup>19</sup> Noteworthy exceptions include Carlsson (2018) and Lewis (2022).

<sup>20</sup> ‘Sufficiently’: I should emphasize here that do not mean to suggest that full intelligibility—either to ourselves or to one another—is either possible or desirable (see n. 26 below; cf. Dover n.d., Gingerich 2022, Gingerich n.d., and especially Russell n.d.). Still, we need to feel the ground is firm enough to, e.g., make joint plans, anticipate one another’s feelings, feel reasonably secure in one another’s affections, and stand ready to help one another in times of crisis. The mistake here—an occupational hazard for philosophers—lies in adopting a maximizing conception of the norm of intelligibility to which these needs give rise. Such a conception makes good sense in many theoretical contexts, as well as in certain practical ones (e.g., in building a bridge or a nuclear plant), but not in one’s relation to oneself or to others.

remain sufficiently attractive (in a broad sense) to others for them to elect to engage in elective relationships with us in the first place. I risk stating the obvious here, but I do it because I think the philosophical significance of this fact has been downplayed, or perhaps even repressed: if we want to have friends, we need at least some people to like us.<sup>21</sup> To have a convenient shorthand to refer to this fundamental feature of social life, I will say that in addition to being intelligible, we also have to be *amiable*—in the eyes of some people—in order to be *eligible* to participate in elective relationships with those people.

Since most of us want and, indeed, need to participate in at least a few elective relationships, most of us want and need to be not only intelligible but amiable to at least a few individual others. It would be surprising if the attempt to become and remain eligible for elective interpersonal relationships had little bearing on our psychic formation—on how our ‘selves’ or ‘identities’ are formed and transformed over time. The pursuit of eligibility is often most acutely palpable (and humiliating) in adolescence.<sup>22</sup> But it would be surprising if our attempts to be eligible for relationships with particular others did not retain the capacity to ‘transform’ us, in the sense invoked by Paul, even well into adulthood.

These two ways in which adults pursue eligibility—by pursuing intelligibility and by pursuing amiability—have not gotten equally respectful treatment from philosophers steeped in the ethos of authenticity. Hill’s treatment of the Deferential Wife illustrates that, to some moral-philosophical noses, the pursuit of amiability reeks of a ‘servile’ disposition to please.<sup>23</sup> The pursuit of intelligibility may look more respectable by contrast, at least at first. Talk of intelligibility, with its implied analogy to language, suggests a relationship to a broader human social world. My being intelligible to you is not merely a private fact of our relationship; it is also a sign that I am, more generally, humanly, folk-psychologically intelligible. Whether you like me is a much more idiosyncratic, contingent matter—and this can make caring about amiability seem, to some philosophers, a bit disreputable.<sup>24</sup>

<sup>21</sup> I use ‘like’ advisedly here, since the word is typically understood to refer to an attitude that no secular philosopher has, to my knowledge, claimed we are morally compelled to have toward everyone. (Even Jesus only demanded that we *love* our enemies; as many readers will know from experience, loving a person does not entail liking them.) By contrast, of course, many philosophical and religious thinkers have claimed that we are morally obligated to ‘love’ or to be ‘friends’ with (at least in some highly moralized and therefore, perhaps, attenuated sense of those terms) all human beings, or at least all of our neighbors or fellow citizens or comrades in struggle. As James H.P. Lewis (2022, p. 4 n.7) has noted, ‘personal liking is...[an] undertheorized notion in moral philosophy’. But Lewis goes on to describe a much more moralized notion than I mean to invoke here, according to which liking sounds a lot like Kantian respect (‘liking a person is a way of treating them as having final, rather than instrumental value’ [Lewis 2022, p. 4]).

<sup>22</sup> The YogiTea oracle addresses a notional customer still painfully wincing at those memories: ‘never try to impress others, try to impress yourself’ (YogiTea 2023g).

<sup>23</sup> Here again, philosophers are swimming with larger cultural tides. In the window of a California yoga studio I recently saw an advertisement for a workshop, entitled ‘Deprogramming People Pleasing: Reclaim Your Power and Take Control of Your Life’, where for fifteen dollars I could learn to ‘live life on [my own] terms’ by ‘set[ting] badass boundaries’ (Loveman 2023).

<sup>24</sup> For an important challenge to this tendency, see Buss 1999. Buss forcefully argues that, far from being a sign of moral weakness, caring deeply what other people think of us—indeed, acknowledging that their conceptions of us ‘[make] a difference to who we are’ (524)—is not only a crucial developmental prerequisite of full-fledged moral consciousness but also a core aspect of adult moral experience. Buss is especially

But the pursuits of amiability and intelligibility are not actually as separate as I have so far made them sound, since the folk-psychological interpretation of another human being is not a normatively neutral endeavor. We all know that hating people tends to make it harder to make basic folk-psychological sense of them, and loving them helps us model their minds. Folk-psychological interpretation and normative evaluation are especially tightly interwoven in the context of elective relationships. In order to form and maintain an elective relationship with you, it is not as though I first try to render you intelligible and then—on the basis of the resultant understanding—decide whether I like you. From the outset, I will be doing both things at once. In order for our relationship to get off the ground, I will have to render you intelligible *as* amiable, and you will have to do the same with me.

This connection between intelligibility and amiability is written into some of the first expressions that English speakers tend to reach for when they try to express affection and devotion. Participants in satisfying relationships often say that they ‘really understand’ or ‘get’ one another. What are they talking about?

They might mean that they are each well-acquainted with the other’s *S* (see Williams, 1981, p. 105). The better I know my way around your *S*, the better I will be at predicting and explaining your behavior. But this cannot be all you mean when you say that I ‘really understand you’. Many, if not most, of our representations of others are normatively valenced. I can dispassionately represent you as having an even or odd number of hairs on your head, but it would be harder to attend for very long to your haircut without having a normative aesthetic response. All the harder to avoid having any normative evaluation of any kind of you as a whole, whether flattering or otherwise. Our overall conceptions of others are ‘thick’: they offer descriptions and express evaluations at once.

An easy way to see this is to try Buddhist lovingkindness meditation. You are supposed to start by bringing to mind someone you love (the ‘Beloved Friend’). You wish the Beloved Friend happiness; that is the easy part. Then you are supposed to do the same with someone about whom your feelings are entirely neutral (the ‘Neutral Person’), and finally with someone you dislike (the ‘Enemy’ or, euphemistically, the ‘Difficult Person’).<sup>25</sup> If we are lucky, it is usually easy to pick a Beloved Friend, and if we are not yet enlightened, it is usually easy to pick an Enemy. But it is notoriously hard to pick a Neutral Person, because it is so hard to come up with someone about whom you really truly have no special liking or disliking.

This exercise illustrates how difficult genuinely normatively neutral representations of other people are to come by. But it also illustrates that our conceptions of others—even people about whom we have very complicated and ambivalent feelings—often have an obvious *overall valence*. In order for an elective relationship to be in decent working order, the overall valence of each party’s conception of the other has to be positive. When I think of a friend, however much anger, resentment, jealousy, or disappointment I might now feel or have felt toward them in the past—and

---

Footnote 24 continued

concerned with the experience of shame. But her point holds more generally: caring about others’ appraisals of us is one of our most basic ways of acknowledging the relevance of evaluative perspectives other than our own (Buss 1999, pp. 532; 537).

<sup>25</sup> The influential Western popularization of Buddhism from which this description is drawn is Salzberg (2008, p. 75ff).

even if, given those complicated feelings, they are hardly the first person I would call to mind as my Beloved Friend for the purpose of what is supposed to be the easy part of the lovingkindness meditation—it ought to be clear that they would, at least most of the time and on the whole, be a better fit for that role than for the Neutral Person or the Enemy. Otherwise, the friendship is seriously, if not yet fatally, endangered.

The fact that you and I stand in an elective relationship means that I voluntarily take up the task of trying to make you intelligible as amiable. The fact that my conception of you is formed in the context of that ongoing task affects its content, as well as the way that content is gathered, structured, and reinforced. To see how, compare Christine Korsgaard's (2009) notion of 'practical identity'. A practical identity is a self-conception that you endorse and treat as a source of reasons, such as 'neighbor' or 'teacher'. In order to play the role Korsgaard assigns it, the valence of a practical identity has to be positive. If you think of yourself as clumsy, that self-conception will causally affect your behavior in all sorts of ways, down to your very gait. But it will not function as a practical identity. A practical identity is not just supposed to help us to understand our lives; it is supposed to guide us through them. To do so, it has to be a self-description that we embrace. Korsgaard argues that we all need practical identities to get us through life. We will be rudderless without some conception of ourselves that we can use both to make sense of our lives retrospectively and to shape them prospectively.

Elective relationships require something similar. For our friendship to work, I need a conception of you that I can actively embrace, allowing it to guide me through my interactions with you. Just as I need a sufficiently flattering self-conception to make sense of my own actions, I need a sufficiently flattering conception of you to make sense of our relationship. Call this my *relational conception* of you. A practical identity is a 'description under which [I] value [my]self and find [my] life worth living and [my] actions worth...undertaking' (Korsgaard, 2009, p. 20, emphasis added). Analogously, my relational conception of you is a description under which I value *you* and find *our relationship* worth preserving and pursuing. My relational conception of you at any given time is the overall picture of you that makes sense, to me, of our relationship, and in particular of the elective character of that relationship—of the fact that I choose to relate to you.<sup>26</sup>

<sup>26</sup> 'Overall': this marks an important disanalogy between Korsgaardian practical identities and relational conceptions. A Korsgaardian practical identity is a *partial* self-description: Korsgaard's examples include things like being a 'human being, a woman or a man, a member of a certain profession, someone's lover or friend, a citizen or an officer of the court, a feminist or an environmentalist' (2009, p. 20). Since each of us has many such identities, we face, according to Korsgaard, 'the task of uniting them into a coherent whole' (21). By contrast, as I am using the term, my 'relational conception' of you is the *overall* idea I have of you at any given time. As this conception expands and evolves, it is subject to norms of folk-psychological intelligibility of the sort that I discuss in the text. But I would want to resist the supposition that we do, or ought to, try to make our relational conceptions of one another 'coherent' or 'unified' in anything like the strong sense that Korsgaard recommends. As the history of the realist novel attests, folk-psychological intelligibility—intelligibility of the sort that makes a fictional character both vivid and plausible—can accommodate much more incoherence, fragmentation, and ambivalence than an ideal Korsgaardian agent is allowed. See Yao (2015) for a forceful critique of norms of 'inner harmony', 'coherence', and 'wholeheartedness' (though Yao focuses on neo-Aristotelian rather than neo-Kantian celebrations of coherence). See also Coates (2023), Gingerich n.d., and Rees (2014).



How does this look from your point of view? Note, first, that you know that I need a relational conception of you in order for our friendship to work. You of course do not use this jargon, but you are aware that, insofar as our friendship is in decent shape, I must have an understanding, image, idea, or working theory of you that is sufficiently rosy to vindicate my devotion to you. You also know that, as the two of us interact and change over the years, I have to maintain and update that conception in light of new events.

This knowledge will invariably shape the way that you behave around me in countless ways.<sup>27</sup> It might prompt you—usually without thinking about it—to play up those of your attitudes and patterns of behavior that you think are central to my relational conception of you. And it might prompt you to downplay or exclude material that might force me to form a picture of you according to which I could no longer value our relationship as I have in the past. You will—again, for the most part unconsciously—endeavor to be, not who you think I *want* you to be (although that might be how Hill would put it), but rather who you think I *take* you to be. Crucially, however, you will not try to be who I take you to be full stop, but who I take you to be *qua* someone with whom it makes sense for me, in particular, to be in this particular relationship.

We can capture all this by saying that, in light of your knowledge that I need a relational conception of you, you construct a *relational identity* for me: a conception of yourself for which I am the notional audience. This notion of ‘relational identity’ can be defined in terms of the more basic notion of ‘relational conception’. Your construction of a relational identity for me is your ongoing attempt to be, for me, someone of whom I can form and maintain a relational conception.

My formation of a relational conception of you is not independent from your formation of a relational identity for me. For instance, my evolving relational conception of you takes into account the fact that you are engaged in constructing an evolving relational identity for me. Who I take you to be, *qua* someone with whom it makes sense for me to be in this relationship, is, in part, someone who wants and tries to be, for me, who I take you to be *qua* someone with whom it makes sense for me to be in this relationship. To double back again, you are implicitly aware of all that—you know that my relational conception of you takes into account the fact that you are forming a relational identity for me, and that you are doing so in the awareness that I am forming a relational conception of you, and so on. In other words, you know that my relational conception of you has been shaped, in significant part, by your attempts to make yourself intelligible as amiable to me. It is because you have had sufficient success in those attempts that I can regard you as an eligible relationship partner, and vice versa.

There is thus an ongoing procedural interplay, as well as a substantive symbiosis, between your attempts to become and remain eligible and my attempts to represent you as such, and vice versa. In the context of an elective relationship that both parties value, each party is trying to understand the other, while also making herself understood: they

---

<sup>27</sup> In their wonderful paper on friendship, Dean Cocking and Jeanette Kennett discuss the ways in which we are ‘drawn, or shaped’ by our friends’ conceptions of us; reflection on this phenomenon, they argue, points toward a conception of the self as ‘a relational thing that is, in part, developed or molded through the friendship’ (Cocking & Kennett 1998, pp. 505–506; cf. Dover, 2022). My discussion here can be read as a further elaboration of some of the mechanisms by which this mutual ‘drawing’ or ‘shaping’ takes place.

are striving for mutual intelligibility. But they want the resultant understanding to be compatible with their mutual endorsement of the relationship: they are striving for mutual amiability. It is their success in these endeavors that makes them eligible to be friends with one another. This brings us back to the everyday expressions I invoked above. It comes naturally to say that two people ‘really understand each other’ when they are each striving for, and so far succeeding at maintaining, mutual eligibility—that is, when each of them is able to make the other intelligible as amiable.

#### 4 How to contain multitudes

In a two-person world, the process I have described would work more smoothly than it does for us. Each of us would have only one relational identity, custom-tailored for the other. In a larger social world, we may have many relational identities—one for each elective relationship that we value. These identities are formed within the context of qualitatively distinct relationships, with a view to these relationships’ being endorsable by qualitatively distinct people. So they differ as inevitably as human individuals do. The more heterogeneous my friends and family are, the more dramatic these differences will tend to be.

These divergences among our relational identities explain why gatherings like birthday parties that bring our friends together can sometimes be a source of acute discomfort.<sup>28</sup> We may feel pressure to switch back and forth among our different relational identities for each guest. But this can be exhausting. So, for the party at least, we may instead just pick one and stick with it. Perhaps we pick the one that we have in the context of the relationship we value most, or the one we have in the context of the relationship we feel most anxious about endangering. But then we may find ourselves uncomfortably aware of the fact that the relational identity we are embodying or enacting will likely be unfamiliar, and perhaps unwelcome in certain ways, to some of our guests.

A more conservative tack would to construct an ad hoc relational identity for the crowd as a whole, which would be a sort of lowest common denominator of our relational identities for each crowd member. That way, we could protect ourselves from the vulnerability that accompanies an awareness that we are acting in ways that some observers might have trouble squaring with their preexisting relational conceptions of us. But while this leveling-down approach guarantees that we will not shock anybody, it also means that we will seem oddly hollow and abstracted to everybody.

We might be tempted to describe this felt tension among our various relational identities in terms of conflict. The conflict framing will be apt if the relevant notion of conflict is understood phenomenologically: at the birthday party, it can indeed be cognitively and emotionally costly to be in the presence of people who have influenced us in very different ways. But the birthday party experience also illustrates that it can be equally costly to try to bring our different relational identities into alignment.

---

<sup>28</sup> The classic sociological study of this phenomenon is Goffman (1959). For discussion of how such tensions can play out across cultural and racial boundaries, see Lugones (1987) and Harris (1990).

An alternative would be to regard our formation of multiple qualitatively distinct relational identities as a case of (at least potentially) benign multiplicity rather than of *prima facie* worrisome intrapsychic conflict. Suppose X and Y are two guests at the party whose conceptions of you are very different; as each of them watches you interacting with the other, they are surprised at what they see. From whose perspective might this scenario seem like a *prima facie* ethical problem? Perhaps from yours: you might worry that the variation means that you are objectionably ‘two-faced’. X and Y might worry, too: they might ask themselves who you *really* are. But such reactions cannot support the conflict framing: they presuppose it. They assume that we ought to strive to be one-faced.<sup>29</sup>

You might instead count our ability to be multi-faced among the riches of human life. You might then replace the conflict framing with a habit of thinking about the differences among your various relational identities as a natural outgrowth of the coexistence of multiple significant interpersonal relationships within a single life, and thus as something to be welcomed. This multiplicity framing would allow you to reflect on your relationships with X and Y and notice how marvelous it is that a single creature is able to create and sustain two such different ways of being. It would not insist on resolving or dissolving the tensions among our relational identities—even when these tensions become uncomfortably salient—by casting some of them as inauthentic or by gerrymandering the ‘self’ to exclude them.

Suppose we accept this latter approach to the differences among our different relational identities. Let me review the assumptions I have so far asked the reader to entertain. I am supposing, with Paul, that the human ‘self’ or ‘identity’, whatever else it is, is a ‘transformable’ sort of thing, and that one way in which it can be transformed is through experience. We noted that encounters with other human beings are central among the experiences that can occasion such transformations. Next, adopting the picture of the pursuit of eligibility described above, I invited you to join me in noticing that one of the ways in which encounters with other human beings can transform us is via our (for the most part, less than fully conscious) endeavors to make ourselves intelligible-as-amiable to them. Moreover, our attempts to make ourselves intelligible-as-amiable to multiple qualitatively distinct others will naturally result in the formation of numerous qualitatively distinct relational identities. Finally, I suggested that we might consider reframing this as a case of benign, or even delightful, multiplicity rather than as a conflict in need of resolution.

Taken together, these assumptions suggest a picture of the self as not only diachronically transformable but also synchronically multiplicitous. Disparate identities, on this picture, do not merely dethrone and succeed each other at dramatic junctures; they cohabit every day, albeit more comfortably on some days than others. On such a picture, worries about authenticity and the concomitant problem of transformative influence would not tend to strike us as quite so pressing or even salient. Interpersonal influence, even transformative influence, is taken for granted on this picture as a normal part of identity formation, and the coexistence of multiple disparate relational identities within a single life is treated as a natural outgrowth of being in a number of different

---

<sup>29</sup> ‘Be consistently constant and constantly consistent’ (YogiTea, 2023h); or, more mysteriously, ‘Be consistent as a constituent part of yourself’ (2023i).

important relationships. Rather than asking ourselves, ‘Which of my relational identities is the true, authentic one?’, we may be more likely to ask a question that I think is much more meaningful (in both senses), namely: ‘Which of my relationships are good ones?’ Adopting this picture of the self will also change the questions we are prompted to ask about the various transformations we notice ourselves undergoing. Rather than asking, ‘Am I too open to influence?’, what if we went back to asking, just like our parents did about us back in the day, ‘Is this person a good influence?’ Rather than asking, ‘Is this what I really want?’ or ‘Will I still be me if I do it?’, what if we simply asked, ‘Should I do this?’ Rather than asking, ‘Am I truly myself around this person?’, what if we instead asked questions like: ‘Do I actually enjoy being around this person? Does this person make me feel free and alive, or dull and constricted?’ (The Deferential Wife might find that last one especially illuminating.)

How we are to answer the latter set of questions is another matter; it may not be much easier than answering the former. And it will, of course, often involve drawing on, and sometimes explicitly reflecting on, the same sorts of psychic contents—‘preferences...dreams, hopes...aspirations’—that the ethos of authenticity encourages ‘deliberating and reflecting on...faithfully modeling...matching [our] choices to...and hewing as close as [we] can to’ (Paul, 2014, p. 105). But these psychic contents will come into play in a very different way when our attention is allowed to roam more freely beyond the self, toward the world and the people with whom we share it.

## 5 Conclusion

Where does this leave us? I have hardly vanquished the appeal of heightened scrutiny. To do so would require loosening the grip of the ethos of authenticity, which would in turn require epochal material and cultural changes—it will not happen in a journal article. (Of course the ethos of authenticity does not entail the heightened scrutiny view: it does not *entail* anything, strictly speaking. But as long as the former ‘lives rent-free in our heads’, as the meme aptly puts it, it will continue to lend plausibility to the latter.) In lieu of a knock-down argument, what I have offered is an alternative way of thinking and talking about the phenomenon of interpersonal influence that does not engage the ethos of authenticity by default. The picture of identity and influence that I have sketched does not entail that we should regard external and internal sources of change with equal scrutiny, but it naturally goes hand in hand with that posture, since it allows us to jump straight to asking substantive questions about our transformations rather than stopping to worry about their provenance.

We have arrived at a choice between two approaches to the phenomenon of interpersonal influence. One draws on the ethos of authenticity and inclines us toward a heightened scrutiny of exogenous influences. The other draws instead on the social aspects of identity and inclines us toward the equal scrutiny of (putatively) endogenous and exogenous influences. Let me close by being more explicit than I have been so far about why I find the latter alternative more attractive.

First, I doubt that the distinction between endogenous and exogenous influences can withstand scrutiny. And even if a general distinction between exogenous and

endogenous influences can be drawn in a principled way, I very much doubt we will be good enough at drawing it *in our own particular cases* to safely appeal to it in evaluating our major life choices, our relationships, and our very identities in the way that the ethos of authenticity encourages us to do. Recall Akhlaghi's Jack and Jill. We are supposed to imagine that Jill has an authentic desire to go to college with which Jack seeks to interfere. But where did Jill's desire to go to college come from? Where do any of our desires come from? It is notoriously hard to say, given how thoroughly human beings are shaped by their social environments. What appear, in adulthood, as our endogenous 'core values' are often just values we absorbed by osmosis in early childhood from the larger culture. By the same token, what appear in adulthood as exogenous incursions may be the activation or expression of previously repressed or otherwise dormant psychic material. In other words, what strikes us as 'endogenous' is often just the psychic stuff that seeps in early, slowly, and osmotically, while what strikes us as 'exogenous' is often just the stuff that hits us faster and harder, after we are old enough to notice.

Skepticism about the celebration of the authentic self is, I think, especially warranted in a monstrously unjust society such as ours, in which so much of what we absorb and imbibe, whether in early childhood or later on, is poisonous. Given what we know about the depth to which human beings are shaped by cultural forces (crudely put, by the influence of large groups of strangers), why should we be any more suspicious of influence that we can trace to a single identifiable individual or handful of individuals—people we actually know, whose character and whose relationships with us we can reflect upon and evaluate on substantive grounds—than we are of psychic forces that seem to come from deep within?<sup>30</sup> All in all, the supposition that we can successfully trace the provenance of our own 'core preferences'—even setting aside the still more dangerously hubristic assumption that the good, to-be-trusted ones are the ones that originate with us as individuals—seems a shaky foundation for an ethics of influence.

Second, I believe that the heightened scrutiny view encourages us to tighten the border controls around our identities to an ethically undesirable extent. I have argued elsewhere (Dover, 2022) that we should think of philosophical theorizing about identity and the self as a first-order normative project, choosing among theories at least in part on ethical and political grounds, rather than just epistemological or metaphysical ones. This is because the notions of the self we embrace—whether these are sophisticated philosophical theories or just 'the culture's ordinary way of thinking' (Paul, 2014, p. 85)—are to a significant extent self-fulfilling. As Iris Murdoch (1999, p. 75) put it, 'man is a creature who makes pictures of himself and then comes to resemble the picture'. I worry that the picture of identity formation that the ethos of authenticity has painted risks alienating us from one another.<sup>31</sup> As Vida Yao (forthcoming, 22)

<sup>30</sup> Diotima's speech in the *Symposium* (Plato, 1997, p. 205e) warns against just this sort of fetishization of what is 'one's own' just because it is one's own, pointing out that 'people are even willing to cut off their own arms and legs if they think they are diseased': we should not understand ourselves as '[taking] joy in what belongs to [us] personally' *as such*, unless by 'belonging to me' we just mean 'good' and by 'belonging to another' we mean 'bad'. Thanks to Ralph Wedgwood for pointing out this connection.

<sup>31</sup> For references to thinkers from a range of philosophical perspectives and traditions who have made this argument, see Dover (2022, p. 218n30).

powerfully argues, there may be ‘pleasures and intensities, intimacies and attachments’ that only become possible if we are willing take the risk of transforming and being transformed by one another. The habit of heightened scrutiny amounts to a way of guarding against precisely this risk—and thereby fending off these forms of intimacy. In short, transformative influence really is a problem for the ethos of authenticity. But one could think ‘so much the worse for transformative influence’ or—as Yao and I do—‘so much the worse for the ethos of authenticity’. Some will want to find ways to reaffirm the ethos of authenticity in the face of this conflict. Others will instead work toward articulating an alternative ethos of relationality. And still others will attempt to find a way to make these two sets of values more compatible than they might at first seem.<sup>32</sup>

Before I conclude, I want to acknowledge an important concern that I hope to address in future work, but to which I cannot do full justice here. The worry runs something like this: ‘relationality’, ‘intimacy’, ‘interdependence’, and other such warm-and-fuzzy concepts sound nice enough at first, but they have a dark side that this paper has so far entirely failed to acknowledge. The mechanisms of influence discussed in §§ 3–4 as indispensable to the formation of elective relationships can also be vectors of terrible psychic harm. Does not embracing an ethos of relationality—with its correspondingly relaxed, non-defensive normative stance toward the psychological realities of influence—leave us more vulnerable to malign and even abusive influences?

While acknowledging the seriousness of this concern, I want to briefly register my conviction that heightened scrutiny is not the right response. True, having relationships leaves us open to having bad relationships. Being influenceable leaves us open to bad influences. But the solution is not to ward off influence by holding ourselves aloof from one another. Many authors (not to mention entire social movements: see Berman [1970] 2009) have prescribed a reaffirmation of the ethos of authenticity as a way of shoring up our self-love and self-knowledge in the face of those who would demean us. But it is a mistake to expect authenticity to swoop in as a norm *ex machina* to rescue us from our descriptively manifest need to be liked and loved.

The case of the Deferential Wife itself affords a convenient illustration of this point. A defender of the ethos of authenticity might counsel her to assert her moral equality with her husband by shoring up her individual identity, standing up for her own needs, desires, and values as she alone can understand and define them.<sup>33</sup> But how many of us can pull off such a feat without first being ensconced among at least a few other people

<sup>32</sup> This third way is pursued by the contributors to an influential anthology (MacKenzie & Stoljar, 2000a) that explores ‘relational autonomy’ as a feminist answer to the ‘hyperbolized ideal of autonomy as self-sufficient individualism’ that dominates ‘the social imaginary of Western societies’ (MacKenzie & Stoljar, 2000b, p. 25). I am in deep sympathy with the criticisms of traditional conceptions of autonomy that motivate these revisionist accounts, as well as with some of the details of the positive proposals offered by the various contributors. But I am much more pessimistic about the prospects for detaching the theoretical notion of ‘autonomy’ from its Kantian roots as a name for ‘the capacity for rational self-legislation [that] is considered to be the defining feature of persons’ (2000b, p. 5). I am accordingly not inclined to return to the philosophical language of ‘autonomy’—let alone to try to rehabilitate the much more culturally redolent notion of ‘authenticity’ (see Meyers 2000)—in the context of an attempt to distinguish liberating from oppressive forms of relationality and community.

<sup>33</sup> ‘May we confine ourselves to the strength of our soul.’ (YogiTea 2023j).

who regard us, not merely as moral equals, but as positively choiceworthy partners in relationships that go far beyond bare moral recognition? It may well be a deficit of such relational sustenance that makes the Deferential Wife's predicament so grim. A big part of her problem, traceable but not fully reducible to the social circumstances that led her to this cretin's bed in the first place, is that she lacks the supra-moral love of particular others whose conception of her she could counterpose to his. Under those conditions, where can a social animal turn for self-love and self-knowledge? What keeps the Wife in her place may not be excessive care for what other people think of her, let alone a 'vice' or 'defect' of 'servility' or 'submissive[ness]'.<sup>34</sup> Her problem could more charitably be reframed as a consequence of a degree of social isolation that no amount of psychic bootstrapping will allow her to escape.

Were the Wife to join a feminist consciousness-raising group, as many previously deferential wives were doing at the time of Hill's article, she might find not only comrades committed to the moral affirmation of her equal personhood, but also—perhaps more essentially—new friends and lovers with whom to make common cause. These new elective relationships might—as they indeed did for many real-life women at the time—afford countervailing influences that might help her get out of her marriage by pulling her in a different direction, toward a different life in which she would have closer friends and better sex and wear more comfortable clothes.

In short, we should at least entertain the possibility that the Wife's problem is not that she is *too susceptible* to influence but rather that she has too few influences—and, so far as we have heard, they are all bad. Reframing her problem in this way would help us to resist the ubiquitous notion that we should not care so much 'what other people think of us'. Indeed, it might help us to get out of the related habit of distinguishing the self from *everybody else all at once*—and apotropaically lumping the latter together as 'other people'—in the first place. It would encourage us instead to recognize, without self-reproach, that there is no solitary way to arrive at a sustainable self-conception. A livable identity requires living among others whose conceptions of us we can unapologetically allow ourselves to care about—even to the point of being transformed by them. As the Deferential Wife case illustrates, this need, in turn, can sharpen our desire for a freer and more livable social world.

I have argued that we would be better off focusing less on figuring out who we really are and what we really want, and more on trying to find wisdom and sustenance in one another. The bad news, from this perspective, is that there is no procedural moral shortcut through the risky, painstaking, substantive lifelong task of making and remaking ourselves as we make and remake our relationships with a finite number of particular, morally flawed 'others'. The good news is that we are not alone in the meantime. To quote the teabag I found most congenial (2023 k): 'love is strength'.

**Acknowledgements** For detailed and thoughtful feedback on drafts, I am very grateful to Farbod Akhlaghi, Jane Friedman, Jonathan Gingerich, Barbara Herman, and two anonymous referees for *Synthese*. I presented versions of this paper at the University of Southern California; at a seminar organized by Ulrike Heuer at University College London; and at the Vrije Universiteit Amsterdam, with comments from Solmu Anttila. In all three cases the discussion was extraordinarily lively and stimulating. Many thanks to all the participants and organizers, and especially to Solmu and Ulrike, for their illuminating responses. The paper also benefited from conversations at a workshop on Transformative Experience organized by Ram Neta and held at Yale

<sup>34</sup> These are Hill's words: see n. 12 above.



University in April 2022, and particularly from exchanges with Vida Yao and others at the session on Yao's 'Eros and Anxiety' (Yao forthcoming). I am also indebted to many other colleagues and friends for discussion and correspondence about these and related issues over the years; I am thinking particularly of Josh Armstrong, Benjamin Bagley, Sarah Beach, Zoe Cullen, James Flynn, Jeffrey Guhin, Pamela Hieronymi, Anthony Laden, Gavin Lawrence, Jessica Moss, Ram Neta, Calvin Normore, Lucy O'Brien, Charles Petersen, Francey Russell, Samuel Scheffler, Seana Shiffrin, Susanna Siegel, Rob Simpson, Amia Srinivasan, Sharon Street, and David Velleman.

## Declarations

**Conflict of interest** I am not aware of any relevant conflicts of interest.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Akhlaghi, F. (2022). Transformative experience and the right to revelatory autonomy. *Analysis*, 20(20), 1–10.
- Berman, M. ([1970] 2009). *The politics of authenticity: Radical individualism and the emergence of modern society*. Verso.
- Boxill, B., & Boxill, J. (2015). Servility and self-respect: An African-American and feminist critique. In M. Timmons & R. N. Johnson (Eds.), *Reason, value, and respect: Kantian themes from the philosophy of Thomas E. Hill Jr.* (pp. 19–41). Oxford University Press.
- Buss, S. (1999). Respect for Persons. *Canadian Journal of Philosophy*, 29(4), 517–550.
- Carlsson, U. (2018). Tragedy and resentment. *Mind*, 127(508), 1169–1191.
- Coates, D. J. (2023). *In praise of ambivalence*. Oxford University Press.
- Cocking, D., & Kennett, J. (1998). Friendship and the self. *Ethics*, 108(3), 502–527.
- Darwall, S. L. (1977). Two kinds of respect. *Ethics*, 88(1), 36–49.
- Dover, D. (2022). The conversational self. *Mind*, 131(521), 193–230.
- Dover, D. (n.d.) *Love's curiosity*. Manuscript on file with author.
- Garnett, M. (2023). Ghostwritten lives: autonomy, deference, and self-authorship. *Ethics*, 133(2), 189–215.
- Gingerich, J. (2022). Spontaneous freedom. *Ethics*, 133(1), 38–71.
- Gingerich, J. (n.d.). *Stop making sense*. Manuscript on file with author.
- Goffman, E. (1959). *The presentation of self in everyday life*. Doubleday Anchor Books.
- Harris, A. P. (1990). Race and essentialism in feminist legal theory. *Stanford Law Review*, 42(3), 581–616.
- Hawkins, J. (2019). Well-being, the self, and radical change. In M. Timmons (Ed.), *Oxford studies in normative ethics* (Vol. 9, pp. 251–270). Oxford University Press.
- Hill, T. E. (1973). Servility and self-respect. *The Monist*, 57(1), 87–104.
- Korsgaard, C. M. (1999). Self-constitution in the ethics of Plato and Kant. *The Journal of Ethics*, 3(1), 1–29.
- Korsgaard, C. M. (2009). *Self-constitution: Agency, identity, and integrity*. Oxford University Press.
- Lewis, J. H. P. (2022). The musicality of speech. *Philosophers' Imprint*, 22(21), 1–16.
- Loveman, K. (2023). *Deprogramming people-pleasing: Reclaim your power & take control of your life*. Advertising flyer posted at Crystalline Collective, Arcata, California. Retrieved July 4.
- Lugones, M. (1987). Playfulness, 'World'-travelling, and loving perception. *Hypatia*, 2(2), 3–19.
- Mackenzie, C., & Stoljar, N. (Eds.). (2000a). *Relational autonomy: Feminist perspectives on autonomy, agency, and the social self*. Oxford University Press.

- Mackenzie, C., & Stoljar, N. (2000b). Introduction: autonomy refigured. In C. Mackenzie & N. Stoljar (Eds.), *Relational autonomy: Feminist perspectives on autonomy, agency, and the social self* (pp. 3–31). Oxford University Press.
- Meyers, D. T. (2000). Intersectional identity and the authentic self?: Opposites attract! In C. Mackenzie & N. Stoljar (Eds.), *Relational autonomy: feminist perspectives on autonomy, agency, and the social self* (pp. 151–180). Oxford University Press.
- Moran, R. (2011). Iris Murdoch and Existentialism. In J. Broackes (Ed.), *Iris Murdoch, Philosopher* (pp. 181–196). Oxford University Press.
- Murdoch, I. (1999). *Existentialists and mystics: Writings on philosophy and literature*. Edited by Peter J. Conradi, with a foreword by George Steiner. Penguin.
- Parfit, D. (1984). *Reasons and persons*. Oxford University Press.
- Paul, L. (2014). *Transformative experience*. Oxford University Press.
- Plato. (1997). Symposium. Translated by Alexander Nehamas and Paul Woodruff. In *Plato: Complete works*, edited by John M. Cooper and Douglas S. Hutchinson (pp. 457–505). Hackett Publishing Company.
- Rees, C. F. (2014). Are intelligible agents square? *Philosophical Explorations*, 17(1), 17–34.
- Russell, F. (n.d.). *Opaque animals: Agency, ethics, and experience*. Manuscript on file with author.
- Salzberg, S. (2008). *Lovingkindness : The revolutionary art of happiness*. With a foreword by Jon Kabat-Zinn. Shambala Publications.
- Starobinski, J. (1988). *Jean-Jacques Rousseau: Transparency and Obstruction*. Translated by Arthur Goldhammer, with an introduction by Robert J. Morrissey. Chicago: University of Chicago Press.
- Strawson, P. (2008). *Freedom and resentment and other essays*. Routledge.
- Taylor, C. (1991). *The ethics of authenticity*. Harvard University Press.
- Trilling, L. (1972). *Sincerity and authenticity*. Harvard University Press.
- Velleman, J. D. (2009). *How we get along*. Cambridge University Press.
- Williams, B. (1981). Internal and external reasons. In *Moral luck: Philosophical papers 1973–1980* (pp. 101–13). Cambridge University Press.
- Williams, B. (1995). Internal reasons and the obscurity of blame. In *Making sense of humanity: And other philosophical papers 1982–1993* (pp. 35–45). Cambridge University Press.
- Williams, B. (2001). Postscript: Some further notes on internal and external reasons. In E. Millgram (Ed.), *Varieties of practical reasoning* (pp. 91–97). MIT Press.
- Williams, B. (2006a). Values, reasons, and the theory of Persuasion. In *Philosophy as a humanistic discipline*, edited and with an introduction by A. W. Moore (pp. 109–118). Princeton University Press.
- Williams, B. (2006b). *Ethics and the limits of philosophy*. With a commentary on the text by A. W. Moore. Abingdon: Routledge.
- Yao, V. (2015). Boredom and the divided mind. *Research Philosophica*, 92(4), 937–957.
- Yao, V. (Forthcoming). Eros and anxiety. *Synthese*.
- YogiTea. (2023a). *Master the unknown by knowing your deep self*. Teabag tag. YOGI TEA GmbH.
- YogiTea. (2023b). *The unknown has to be trusted*. YOGI TEA GmbH.
- YogiTea. (2023c). *You should develop mental strength*. YOGI TEA GmbH.
- YogiTea. (2023d). *Get in touch with yourself*. YOGI TEA GmbH.
- YogiTea. (2023e). *Your deep inner self is your brightest light*. YOGI TEA GmbH.
- YogiTea. (2023f). *Without realising who you are, happiness cannot come to you*. YOGI TEA GmbH.
- YogiTea. (2023g). *Never try to impress others, try to impress yourself*. YOGI TEA GmbH.
- YogiTea. (2023h). *Be consistently constant and constantly consistent*. YOGI TEA GmbH.
- YogiTea. (2023i). *Be consistent as a constituent part of yourself*. YOGI TEA GmbH.
- YogiTea. (2023j). *May we confine ourselves to the strength of our soul*. YOGI TEA GmbH.
- YogiTea. (2023k). *Love is strength and will give us knowledge*. YOGI TEA GmbH.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.