

Epistemic Self-Alienation

Vlad-Stefan Harbuz

It was always clear to Sophie what she wanted to do when she grew up — she wanted to be a tailor. As a teenager, she would often stay up until sunrise sewing clothes for her friends, and her creations won awards in local competitions. While her passion never waned, she found herself dedicating less and less time to sewing as the moment to decide on a career neared. She figured that most successful people go to university, and was encouraged by her family to get a degree in economics — it pays well, and after all, she didn't hate working with numbers. The path before her lay itself almost automatically. After five years of studying, her university's careers service encouraged her to find a job and guided her to a well-paying position at a large consulting company. This position required her to move to a big city in which the inflated house prices would always require her to have a high salary in order to survive, but after all, this is what adult life is like. Her life was now “on track”. Six years later, Sophie is 30, and she has spent most of the waking hours of her adult life optimising balance sheets to bring value to shareholders.

Many individual decisions came together to lead her to this stage in her life, but what is curious is the manner in which these decisions were made. Sophie never cared about economics or consulting companies. The decisions leading her to pursue these things do not align with what is important to her, so there is a sense in which we can say that these decisions are not her own. They happened automatically, guided by a generic inherited notion of what one “should do”. Her family praises her achievements and prestigious role, but this very act feels unreal to her, as her achievements feel like they are not her own. When her friends ask her why she chose not to become a tailor, she says something about financial stability. This makes sense at first — one has to earn money after all. But there is no shortage of tailoring jobs in her country, and the income earned in these jobs would be more than sufficient, because she has no expensive interests or tastes. One might think that she has simply lost interest in her previous passion, but this is not so — every time she sees a finely-constructed garment, she cannot contain her excitement. She could not stop caring about sewing even if she tried to. Nevertheless, she has ended up in a career that her young self would find unthinkable. No prudential considerations stopped her from becoming a tailor, but in her inventory of possible decisions, it was somehow never a real option. Her life seems to not be her own, and as time goes on, it

feels more and more wrong to her.¹

In this essay, I would like to explore the question of which framework would be most suitable for characterising Sophie's situation. In order to describe her example, I will introduce the concept of *epistemic self-alienation*, a diagnosis which stands at the intersection of self-alienation and epistemic injustice. In order to define epistemic self-alienation, I will consider two general questions. Firstly, what is *self-alienation* really, and what does it mean to say that Sophie's life is not her own? The most general definition of alienation describes it as a problematic separation of a subject and an object that belong together, with the subject usually being a person (Leopold, 2022). In the case of *self-alienation*, the subject and the object are identical, meaning that the person is alienated from themselves. But being problematically separated from oneself sounds confusing at first. What parts of oneself might one be problematically separated from? And what aspect of this separation makes it problematic? Even though I am not attempting to give a complete definition of self-alienation, but rather only of a special case of it, namely epistemic self-alienation, we should still have some kind of understanding of what self-alienation actually is, through the lens of existing literature. I will use Harry Frankfurt's conception of "caring" in order to suggest such a definition, then test it against the examples of self-alienation proposed by Jaeggi.

Secondly, what does it mean for Sophie to be *epistemically* self-alienated in particular? It is clear that Sophie was somehow limited in her capacity to make certain decisions — certain courses of action were simply not available to her. I will describe this idea in more detail and, in order to explain why it is an injustice, I will show that it is closely related to Miranda Fricker's idea of epistemic injustice, specifically hermeneutical injustice.

In short, in this paper I claim that we are epistemically self-alienated to the extent that we are systemically rendered unaware of the possibility of prioritising the things we care about, which inflicts an epistemic injustice upon us.

I believe that most people would consider a life affected by this injustice less desirable than one in which we prioritise things that are important to us. At the same time, deprioritising one's intrinsic reasons and instead prioritising a set of socially inherited extrinsic reasons is a widespread phenomenon. It has been described by Marx in his definition of "ideol-

¹I am grateful to Rahel Jaeggi for formulations such as "one's life being 'on track'" and "one's life not being one's own" (Jaeggi, 2014, p. 52).

ogy” (Marx and Engels, 1965), and further expanded in, for example, Herbert Marcuse’s description of “false consciousness” (Marcuse, 2013). Because this phenomenon is both harmful and widespread, I believe the concept of self-alienation in general, and epistemic self-alienation in particular, is one of the most important ideas to discuss when it comes to deciding how we should live in our current time.

1 Value and Caring

First of all, I would like to define some properties which will be essential to us, namely instrumental value, intrinsic value and caring. We can consider something instrumentally or intrinsically valuable, and we can care about something. These are all relational (or extrinsic) properties, in that, if I care about tennis, it is not that tennis has the property of being cared for, but rather, I care for tennis in virtue of my relationship to it (Lewis, 1983).

1.1 Instrumental Value

If we value a certain object instrumentally, we consider it desirable as a means to an end — this object’s value lies only in the fact that we believe it will bring about some other desirable thing ². There is something notable here: an object we value instrumentally is never so valued by itself, but always in combination with some other object. Therefore, a 1-to-1 relation is essential for instrumental value. This allows us to imagine valuing A instrumentally in order to bring about B, which we value instrumentally in order to bring about C, forming a kind of chain — I will describe this in much more detail later.

Another thing to note is that instrumental value can be fungible ³. This means that an object we value instrumentally might be replaced by any other object providing equivalent value. Andreas Wildt says that “what (...) is a mere means is something to which the agent is indifferent in a certain respect; it could be replaced by any other equally suitable means without any loss” (Wildt, 1987; further discussion in Jaeggi, 2014, p. 208).

For example, let’s say I instrumentally value going to work, because it allows me to satisfy the end of earning money. This could just as easily be replaced by another equivalent instrumentally valued object. If playing tennis also allows me to earn money, I lose nothing

²My definition is based on Max Weber’s (Weber, 1978, p. 24), but I find his definition less clear.

³Something is fungible if it can be replaced with another equivalent item.

if I play tennis instead of going to work, at least all other things being equal. There is no value in this example that going to work has but playing tennis lacks, or vice versa.

There is one last thing I would like to note, namely that instrumental value seems to be *impersonal*. If going to work allows me to earn money, this gives me a perceived reason to go to work if I want to earn money. But this reasoning does not apply only to me in particular. If my reasoning is sound, and someone else wants to earn money, this also gives them a perceived reason to go to work. To be sure, my reasoning can be right or wrong, but if it is right, it does not apply to me in virtue of some particular quality of mine.

1.2 Intrinsic Value

If we value something intrinsically, we value it for its own sake — our valuing it is not dependent on connecting it to some further end. Unlike instrumental value, intrinsic value is personal. Us finding intrinsic value in some object does not provide grounds for someone else to find intrinsic value in it.

Intrinsic value, like instrumental value, can also be fungible. Harry Frankfurt provides the example of a person faced with the decision between going to a concert or spending their time with a friend. This person considers both activities intrinsically valuable, but ultimately, this person's interest lies in spending an evening in *some* intrinsically valuable way. Therefore, “a person who desires something exclusively for its intrinsic value (e.g., listening to music) may be happy to accept instead something else (e.g., helping a friend) that also possesses intrinsic value but possesses it by virtue of characteristics of a quite different sort” (Frankfurt, 1999, p. 164).

1.3 Caring

Some things that we find intrinsically valuable happen to also be things we care about. One might consider things such as cooking or shoemaking to be worth doing for their own sakes without personally caring about these things at all. Harry Frankfurt says that caring about something entails considering it important to ourselves (*ibid.*, p. 155), but we should want to figure out what this means more specifically. I believe there are at least four important aspects that will help us define what it means to care about something.

Firstly, caring about something means that we will *necessarily* be harmed if we are sepa-

rated from it. Suppose that I really want to attend a concert, but it turns out that I cannot go. I might not be able to go for a very good reason, such as meeting a good friend, and I might even prefer not attending the concert to attending it under these circumstances. But even so, if I care about going to the concert, my separation from it will necessarily hurt. Our connection to a cared-for object cannot be replaced by a connection to some other cared-for object. Because of this, caring is non-fungible. If I care about both playing tennis and painting landscapes, it could hardly be said that I should not be injured by my inability to paint simply because I am able to play tennis. This would be akin to saying that one should not feel a sense of loss in being apart from one's partner simply because one is able to spend time with one's parents.

Secondly, care exhibits a kind of persistence — my desire to attend the concert persists long after I have decided not to attend (Frankfurt, 1999, p. 159). Still, it cannot simply be that the persistence of my desire to attend the concert happened by accident or through inertia (ibid., p. 160). I must *want* to desire attending the concert. In Frankfurt's terms, caring does not only involve a first-order volition, for example the desire to attend the concert, but also a second-order volition, the wish to continue to desire to attend the concert. My practical arrangement with my friend might override my first-order volition to attend the concert, causing me to decide not to attend, but it cannot override my second-order desire "that this first-order desire not be extinguished or abandoned" (ibid., p. 161). If separation from X necessarily causes us harm, then (1) we *need* X, (2) we have an interest in *wanting* X, but also (3) we have an interest in *wanting to want* X.

Thirdly, an essential component of caring is also that we often cannot help what we care about. If someone asked one why one cared about something, one might not be able to answer anything more than "I just do", in the same way that one could not justify one's love for one's partner. In fact, we may wish that we did not care about something, and we may regard caring about it as irrational, but our caring might persist nonetheless.

Lastly, and perhaps most importantly, we can now see, based on the above, that caring about something gives us an altogether more special connection with it than simply finding it valuable. For Frankfurt, because caring is a bond to something from which separation will necessarily cause us harm, caring is also an important looking glass into the objects that we have a strong and insurmountable *need* for. This also means that caring is essential to our identities. "Our essential natures as individuals are constituted

(...) by what we cannot help caring about” (Frankfurt, 1999, p. 138).

There are two readily apparent objections to this conception of caring. First of all, does separation from a cared-for object really *necessarily* cause harm? I could care about my partner, but my relationship to my partner might be abusive, meaning that the relationship causes me much more harm than the separation would. This is, however, not a problem for our formulation. It may certainly be true that separation prevents more harm than it causes, but in this situation too, separation necessarily causes us some harm, even if it is lesser than some other harm caused by non-separation.

Secondly, there is an implication in my above definition that the harm caused by separation is an injustice. But is it really an injustice for me to be separated from my abusive partner, even if this causes me harm? I claim that separation from a cared-for object, and the ensuing harm, is an injustice *by default*. However, this possibility for injustice can be overridden and trumped by other considerations, for example, the fact that the partner I am being separated from is abusive. If such a justification is present, the separation is not an injustice. However, if a sufficient justification cannot be identified, we should default to seeing this harmful separation as an injustice.

The above concepts will be useful throughout this paper, and now allow us to turn back to Sophie’s situation.

2 Epistemic Self-Alienation

Let us now try to understand what is problematic in Sophie’s case. Conveniently enough, in chapter 5 of her book “Alienation”, Rahel Jaeggi describes an example that is strikingly similar to Sophie’s (Jaeggi, 2014, p. 52). Fortunately, this means that Jaeggi’s analysis of her example is also useful in our case. I believe that a good start to describing what is wrong in Sophie’s case is that her life becomes, in Jaeggi’s words, “rigidified” (ibid., p. 54).

What does it mean for someone’s life to become rigidified? This term describes a life in which decisions are automatically “made” in advance, before the question of what should be done is even posed. “Everything appears unalterable”, as if it consisted of “fixed, invariable components that one no longer has access to”. This is a corruption of the situation we really ought to find ourselves in, namely one in which we “[realize]

that the relations in which one lives are in principle malleable but also (...) [come] to see that decisions already made are fluid and open to revision” (Jaeggi, 2014, p. 59). This matches descriptions of reification or naturalisation (Lukács, 1971) — “something made turns into something given and outside one’s command; the agent’s own actions (or their results) confront her as an alien power” (Jaeggi, 2014, p. 57).

But such rigidification is not, by itself, *necessarily* harmful. There might be aspects of my life that I come to relate to in such a way such that they appear unalterable, but that after all are not important to me at all. Say my choice of electricity provider becomes rigidified, but this choice is ultimately meaningless to me to begin with. Nothing significant is lost, and to say that my rigidification is an injustice seems excessive. Jaeggi certainly makes a powerful point when she says that “practical questions must be posed not just once but over and over, even with respect to familiar, longstanding practices” (ibid., p. 60), but is this really true of all questions?

In fact, only a subset of cases of rigidification are problematic. We can understand these cases better by noting some very interesting phrasing that Jaeggi uses in passing. She says that rigidification occurs when “a process **that really ought to be a result of actions** — or that by its nature is an action process — appears (or must appear) to the participants as taking place independently of human agency” (emphasis mine, ibid., p. 60).

It therefore seems that a clearer definition of what is problematic in Sophie’s case must have at least two components. Firstly, rigidification must be present, the result of which is a kind of “masking of practical questions” (ibid., p. 52), an epistemic blindness in which our very capacity to decide becomes opaque to us. This is all clearly apparent in Sophie’s example. So many of Sophie’s decisions have happened almost automatically: it seemed to just “make sense” for her to choose a career in economics, and once she did so, it similarly made sense for her to get a job at a consulting company and move to a metropolis where such jobs are plentiful. Her passion for sewing has been fully neglected, not because she ever purposefully set out to neglect it, but through a kind of inertia exerted upon her by the life path she has defaulted to, rendering her blind to the fact that she could even actually choose to become a tailor.

Jaeggi's commentary on her own example describes this phenomenon eloquently:

This does not merely mean that she has not acted, or has not availed herself of her possibilities for acting, **but that she has not even understood her situation as one in which action is called for or possible**; it does not merely mean that she has not decided something for herself, or has not led her life herself, **but that she has been incapable of understanding or regarding it as something she can or must lead**. (Jaeggi, 2014, p. 56) (...) If her situation is to be understood as reified, this is not because the practical question is answered by someone other than herself; rather, it is not even posed. ⁴ (ibid., p. 201)

Secondly, it seems essential that these decisions “really ought to be a result of actions”. That is, it is problematic for us not to be able to participate as an agent in a certain decision if the matter has some kind of importance to us.

But what things “really ought to be a result of actions”, and what makes these things “important”? In which ways must Sophie value or relate to her life path or to sewing for her separation from these things to be problematic? This is where the conceptions of value we described earlier prove useful. Let us consider these conceptions in turn, starting with instrumental value. It does not seem suitable to say that it is problematic to not be able to appropriate a decision to do something instrumentally valuable as our own. After all, there are many instrumentally valuable objects that we could prioritise in our lives, but their very nature seems to imply that they are not important to how we constitute our identity and relate to ourselves, as they are only a means to a different end.

We might think that being problematically separated from objects we value intrinsically could constitute an injustice. But as we have seen previously, there is a whole host of possible intrinsically valuable things that we have no special connection to. Perhaps I find painting intrinsically valuable, but I personally have no specific connection to painting, and no particular desire to appropriate this activity as part of my life. It could hardly be said that an injustice is inflicted upon me if some kind of epistemic blindness keeps me from the possibility of painting — after all, it's all the same to me either way.

What does, however, seem to capture what we are looking for, is Harry Frankfurt's notion of caring. The suitability of this notion is suggested by the fact that caring entails that

⁴I have edited this quote to change the pronouns, and the emphasis is mine.

we have a commitment to preserving our relationship to some thing, because we consider apprehending it personally important to us. I therefore suggest that, in order to be able to say that our rigidification in relationship to something is an injustice, we must care about that thing.

This defective relationship to objects of care is exactly what I believe is problematic in Sophie's situation. Sophie is rendered epistemically blind to the possibility that she could appropriate for herself aspects of her life such as her education, her career, and the place in which she lives, even though there is no prudential or instrumental barrier to doing so. Though Sophie's passion is sewing, she is blind to the possibility of prioritising this aspect of her life even when sewing jobs are plentiful in the place in which she already lives. It never even occurred to her that she could do anything other than follow the predetermined career path that was automatically laid before her. This is an injustice because Sophie's life path is something that she should have agency over, because it is something she cares about. It is this epistemic inability to relate to cared-for objects that makes Sophie's case particularly problematic.

Based on this, I propose the following definition:

Definition 1 *One is epistemically self-alienated to the extent to which one is epistemically blind to the possibility of prioritising things that one cares about.*

While it's certainly encouraging that we've reached a definition that seems to bring together our discussion so far, we must evaluate this definition to see if it holds up to scrutiny.

2.1 Caring as a Criterion

I would like to start off by examining whether caring is really a necessary condition of epistemic self-alienation. If it is, two things should be true. Firstly, if we do not care about something, an injustice cannot be inflicted upon us as a result of our relationship to it. This seems true because, as we have seen in the above examples of instrumentally valuable and intrinsically valuable things, if we do not care about something, there is an upper bound of harm that separation from it can cause.

Secondly, it should be true that, assuming that the other conditions for epistemic self-alienation are met (e.g. rigidification), an injustice is inflicted upon us if we are rendered

unable to appropriate a cared-for thing. This seems to arise naturally from the very definition of care. A cared-for object is one that we relate to in such a way that separation from it will necessarily cause us harm. But being blocked from positively relating to cared-for objects seems to mean that we are rendered *unable to avoid harm*, which is surely an injustice, at least in the default case. As I have described in the above section on caring, while being separated from a cared-for object is not always an injustice, it is certainly an injustice if there is no justification for this separation and the harm it causes.

A further reason is that things we care for form the basis of our identities. Kambartel claims that “we realize ourselves in actions that we do for their own sake, on the basis of ‘our own’ decision, namely, on the basis of a decision we insightfully make in relation to our life” (Kambartel, 1989, p. 24; further discussion in Jaeggi, 2014, p. 206)⁵. Frankfurt goes further and claims that caring is important to our identities independently of the cared-for object:

The significance to us of caring is thus more basic than the importance to us of what we care about. (...) Caring is important to us for its own sake, insofar as it is the indispensably foundational activity through which we provide continuity and coherence to our volitional lives. Regardless of whether its objects are appropriate, our caring about things possesses for us an inherent value by virtue of its essential role in making us the distinctive kind of creatures that we are. (Frankfurt, 1999, p. 162)

It then seems that our relationship to cared-for things is absolutely something that “really ought to be a result of actions”, and that having a defective relationship to cared-for things causes us to have a defective relationship to ourselves. This seems to suggest that our relationship to cared-for things is a good criterion for whether or not we are epistemically self-alienated.

2.2 A Subtype of Self-Alienation

Another question that is worth asking is whether the concept I am describing warrants the name of “self-alienation”. To thoroughly examine this, we would have to find a good definition of self-alienation in general, which is significantly beyond the scope of this

⁵In this quote, Kambartel is referring to intrinsically valuable things, not cared-for things. But cared-for things are a subset of intrinsically valuable things, so this quote necessarily applies to cared-for things as well

paper. However, I think it is worth consulting what Rahel Jaeggi says about alienation to determine whether our definition is compatible with her formulation.

How does Jaeggi, then, define alienation? The following passage is the clearest definition I have found in her work:

Someone is alienated, I have claimed, if she cannot react to her own given conditions. Since the sociality of our existence belongs among these conditions, the alternative between freedom and alienation is decided by how and to what extent we succeed in making this sociality our own. (Jaeggi, 2014, p. 219)

This definition is a good start, but we must further clarify it. What does it mean to “make something our own” in this sense? Clifford Geertz says that “the question (...) ‘How Comes It that we all start out Originals and end up Copies?’ (...) finds (...) an answer that is surprisingly reassuring: it is the copying that originates” (White and Epston, 1990, p. 13; referencing Geertz, 1986, p. 380). We might interpret making socially inherited ideas our own as meaning that we adopt them in a way that necessarily sees us modifying them to suit our own needs, as opposed to taking them on as an unmodified monolith. White and Epston claim that “the evolution of lives is akin to the process of reauthoring, the process of persons’ entering into stories, taking them over and making them their own” (White and Epston, 1990, p. 13).

Of course, whether one can make one’s sociality one’s own also greatly affects the outcomes of one’s life. White and Epston refer to the set of stories that are “basic and common to the members of a social category although occurring independently to each of them” as the “dominant story”, while “aspects of lived experience that fall outside the dominant story” are termed “unique outcomes” (ibid., p. 15). It is exactly this dominant story that embodies a Foucauldian subjugation — adherence to the dominant story and separation from unique outcomes “forges persons as ‘docile bodies’” (ibid., p. 20). Indeed, Foucault saw this conformity-enforcing “everpresent gaze” as having “replaced the judiciary and torture as a primary mechanism of social control” (ibid., p. 24; referencing Foucault and Mailänder, 1975).

To make this as clear as possible, let us consider an example of such as story, or, as White and Epston call it, a “narrative”. Sophie might inherit a narrative that looks something

like “sewing is not a *real* job”. The way in which Sophie engages with this narrative is likely to impact how she ends up living her life — “the stories or narratives that persons live through determine their interaction and organization [and] the evolution of lives and relationships occurs through the performance of such stories or narratives” (White and Epston, 1990, p. 12).

If Sophie simply adopts the dominant story that “sewing is not a real job”, she experiences an “[obliteration of] the opposition between the private and the public existence, between individual and social needs” (Marcuse, 2013, p. xlvi). For Sophie to have a unique outcome, to make this narrative her own, she must question and “retell” it in some way. To put it simply, Sophie must decide if *she* thinks sewing is a “real” job. “As persons become separated from [the dominant story], they are able to experience a sense of personal agency (...) a capacity to intervene in their own lives and relationships” (White and Epston, 1990, p. 16).

These descriptions of what is problematic about self-alienation seem to match Sophie’s case quite well. She could have reauthored and retold her story, made her story her own, and found a unique outcome by rejecting the epistemic oppression that compels her to work for a consulting company out of inertia. But instead, Sophie has unfortunately failed to reauthor her story, and has fallen prey to the dominant story, which says that sewing is not a real job. She has failed to react to her pre-given conditions, and she has not made her sociality her own, at least in this respect. Therefore, while I cannot give a full account of self-alienation in general, it seems that our definition of a particular subtype of self-alienation, namely epistemic self-alienation, is very much in line with Jaeggi’s contemporary conception of self-alienation.

2.3 Substantive vs. Procedural Definitions

Another strength of our definition is that it is a *procedural*, as opposed to *substantive* definition. I will explain what these terms mean, why they are important, and why our definition is procedural.

In short, a substantive definition focuses on whether or not it is appropriate to value an object, whereas procedural definitions will focus on whether the process by which we came to value something is valid, without passing any judgement of what we should and should not value.

Substantive definitions imply that we each have some kind of essence (or, in Marxist terms, species-being) that we can either have a positive connection with, or be alienated from. Jaeggi describes such a model of alienation as a “core model” which claims that “one is ‘with oneself’ (bei sich) or authentic when one is in agreement with an inner essence or with a kind of internal ‘original pattern’ of oneself”. If we find ourselves separated from this essence, our only hope is to rediscover it and “find ourselves” (Jaeggi, 2014, p. 157).

These essentialist conceptions of alienation have many problems, which should lead us to avoid them. Firstly, we are faced with the apparent paradox of how “something that does not correspond to its own essence simultaneously partakes of it and does not partake of it” (ibid., p. 45). Secondly, it seems natural to ask how we might hope to find this “core” if it is so separate from our external manifestations.

But most importantly, such definitions can be paternalistic and reductive. Any substantive definition of alienation will deal with whether it is appropriate to value an object, but universally decreeing what should and should not be valued, no matter at how general a level, cannot hope to capture our individuality as people ⁶. Therefore, we should hope to describe alienation without such a core model.

It is easy to see why a procedural definition is, in many ways, preferable to an essentialist one. Because we do not have to resort to finding an essence, we can examine the totality of one’s actions, which constitute one’s self, and say if there is anything defective about them. Because one’s self is necessarily created through expression, it is also more straightforward to examine one’s so-construed self than one’s “essence”. And the procedural nature of this definition means that we diagnose self-alienation by saying that someone has related to themselves in a deficient way, without prescribing what their core should be. In Jaeggi’s words:

If processes of self-alienation can be analyzed as ways in which actions are disrupted or constrained rather than as a falling away from or a distortion of a substantial essence, then one is not alienated from something but in the performance of an action. (ibid., p. 160)

⁶Natalie Stoljar makes some interesting counterarguments to this (Stoljar, 2000). I have tried to structure my arguments so that they are not vulnerable to Stoljar’s objections, but I will not discuss this for reasons of space.

To this end, Jaeggi describes a model in which “the true self is no longer an ontologically independent entity”, but rather “the sum of one’s actions”, which is always undergoing change and development (Jaeggi, 2014, p. 164). It is then “not possible to make out any ‘doer behind the deed’ (Nietzsche), no self that exists prior to and apart from the deed” (ibid., p. 160). To reinforce this point, she quotes Hegel as saying that “an individual cannot know what he is until he has made himself a reality through action”, and claims that “we cannot be separated from how we express ourselves in the world (in what we do and say)” (ibid., p. 165).

Fortunately, we have reached exactly this kind of procedural definition. We have said the following:

Definition 1 *One is epistemically self-alienated to the extent to which one is epistemically blind to the possibility of prioritising things that one cares about.*

At no point are we prescribing what Sophie *should* care about. Rather, we are simply saying that Sophie does care about certain things, and that her caring for these things can be seen in her actions, in how she expresses herself in the world, for example when she cannot help but feel hurt when she is separated from something that she cares about, such as sewing. In fact, we are sticking quite close to Jaeggi’s description, because our epistemic blindness is exactly the kind of way in which “actions are disrupted or constrained”, whereas we do not ever refer to the “distortion of a substantial essence”. Therefore, our definition is a procedural and not substantive one, which means that it does not suffer from the weaknesses of substantive definitions, which is a point in our definition’s favour.

2.4 Two Objections

We have seen that our definition handles Sophie’s example quite well. At the same time, this is no surprise — it is, after all, our paradigmatic case. One might object that our concept of epistemic self-alienation is overly specialised to only a few specific cases of self-alienation, and that Sophie’s example merely conveniently fits. However, I would like to show that this idea of epistemic self-alienation actually applies much more broadly than it appears to at first sight. To do this, I will show how our definition can be useful in understanding one of the examples of self-alienation Jaeggi gives in her book. In fact, I claim that our definition gives us a richer perspective on this case than Jaeggi manages to.

The example can be found in chapter 7 of “Alienation”:

H., a self-professed, reflective feminist of strong convictions, catches herself over and over again communicating with her lover like a silly, giggling adolescent girl. She rejects such forms of feminine coquetry as unemancipated, as the mannerisms of a “little girl”. She has long understood that the idea that women must present themselves as cute, petite, and harmless in order to be attractive is the projection of a world dominated by men. Yet, as she discovers to her irritation, she constantly falls back into these patterns of behavior against her will. She experiences her own behavior, so starkly in contrast to her convictions, her self-conception, and her life plan, as contradictory and as not really part of herself. It triggers in her a feeling of disconcertedness when she sees herself behave in such a manner: “That can’t be me”. Formulated somewhat dramatically, it is as if in her giggling something were speaking through her that is not herself. (Jaeggi, 2014, p. 100)

We should want to know how Jaeggi characterises this case of self-alienation. She says that, in relation to this example, “self-alienation means not being able to move freely in one’s life, being inaccessible to oneself in what one wants and does” (ibid., p. 128). What Jaeggi seems to find problematic about this case is that there is some kind of internal division or conflict between various aspects of H.’s actions. Her desire to act like a giggling girl inhibits her appropriation of her feminist ideals, and her feminist ideals block her from communicating with her lover in a silly way, which is how she sometimes wants to express herself. “H.’s internal division hinders her in doing what she really wants, in being able to move freely in her life” (ibid., p. 130).


This is all well and good, but it does not clarify how and why it is possible for H. to simultaneously view a desire as part of herself and also alien from herself. It also does not suggest what a resolution might look like, which is unfortunate, because it seems useful to know what the equivalent non-self-alienated situation would be. Must H. decide that either the giggling girl or the feminist is who she “really” is, and discard the other? Can the two be reconciled? Or is there some other solution?

I would like to talk about how the idea of epistemic self-alienation might apply here. At first sight, our definition does not seem to apply to this situation. Is H. really self-alienated to the extent to which she is epistemically blind to the possibility of prioritising

things that she cares about? It is straightforward to see how *Sophie* can suffer from an epistemic blindness that stops her from prioritising things she cares about. But this does not seem to apply in H.'s case, because H. seems to care about both acting like a giggling girl *and* being a reflective feminist — the things she cares about are in conflict with each other. In addition, she does not seem to be epistemically blind, as she is conscious of both parts of her personality — both the possibility of acting as a reflective feminist and that of acting like a giggling girl are visible to her. This seems to provide an argument against our definition.

However, I think that our definition actually applies quite well. First of all, I believe there *is* actually an epistemic blindness happening in H.'s case. It is clear that she cares for, but also rejects, being a giggling girl, and the same is true for being a reflective feminist. It seems impossible to reconcile the two. But to frame the situation as one in which H. needs to choose one aspect of her personality or another as being authentic is to view the situation superficially. Jaeggi herself acknowledges that, if we are to view H.'s options as a dichotomy, she must choose between dedicating herself to living as a giggling girl or reflective feminist, but “both might be cases in which being in agreement with oneself is achieved at the price of a loss of self and where her self-conception as a whole would be false, manipulated, inappropriate, or illusory” (Jaeggi, 2014, p. 118).

It therefore must be that H.'s view is insufficiently differentiated. By this I mean that she views the archetype of the giggling girl, and the archetype of the feminist, as monoliths, when in fact there are components of both that she authentically cares for, and components of both that she rejects. This is what I mean by “insufficiently differentiated” — H. is not observing the higher-resolution component parts of these perceived monolithic attitudes. Ideally, H. would construct an identity that is not limited by these two arbitrary categories, and that incorporates aspects of both. However, it seems that H. is constrained to think only in terms of these two concepts: the giggling girl and the reflective feminist. In this respect, I believe she suffers from an epistemic blindness.



In Jaeggi’s example, H. is only conceptualising things in terms of “giggling girl” or “reflective feminist”, but what I have labelled “H.’s cares” above could be conceptualised as a category in and of itself, albeit incorporating elements from each of the other categories. In H.’s example, this category instantiates what Theodor Adorno referred to as “the non-identical” (Adorno, 1973), that which belongs to a conceptual category which we cannot perceive, and we are therefore implicitly (and regrettably) compelled to erase.

Therefore, I suggest that our definition stands up to the test of H.’s example, because this example, too, can be interpreted as one in which H. is epistemically self-alienated, as she is epistemically blind to the possibility of prioritising the things she cares about. Specifically, she is blind to the differentiated category of her cares, which she is therefore unable to prioritise.

There is one other potential objection I would like to discuss. Naturally, the epistemic self-alienation we have been discussing originates from social factors — inherited ideas, peer pressure, fear of poverty, the demands of the market, and so on. There is an obvious but faulty objection that might easily arise, namely that our diagnosis of epistemic self-alienation is excessive because this supposedly self-alienated behaviour is inevitable in a market-based society. More generally, one might object that my definition of epistemic self-alienation says that it is an injustice for people to not be able to do what they want, which would seem like a silly definition. What might we make of this?

Waheed Hussain describes markets as dynamical systems formed by the “production activities and consumption activities” of agents. These systems “constantly [adjust] and [readjust] to changing circumstances in order to maintain a pattern that is economically efficient” (Hussain, 2023, p. 10). This also necessarily means that markets limit the op-

tions that are available to us — if consumer preferences change, the market might lead me to the healthcare sector by scaling back my options in car manufacturing, thereby providing an incentive which ends up changing and limiting my ultimate course of action (Hussain, 2023, p. 6). But this means that, if we are to achieve efficiency, not all options can be open to everyone. It is in this way that one might claim that my conception of self-alienation does not make sense, because my claim ends up looking like I’m suggesting that it is an injustice for individuals to not be able to do what they want, even if these wants are detrimental to society and grossly conflict with the market.

Of course, these criticisms can be pushed back against using an immanent critique of markets — Hussain himself points out that “these changing option sets draw individuals into patterns in a way that bypasses their private judgements about the merits of these patterns” (ibid., p. 10), hence limiting the agents’ freedom and failing to meet an anti-authoritarian ideal. But my point is entirely different, namely that our discussion of alienation entirely bypasses discussions about the efficiency of markets.

Sophie is faced with at least two options, becoming a tailor and working for a consulting company. But, importantly, and as we saw in the introduction, both of these options are similarly efficient from a market perspective. Both services that Sophie could offer are in demand, and jobs are available for both of them. Therefore, our question is: why is it that social factors can drive Sophie away from her cared-for projects, and towards projects that are merely instrumentally valuable, *even when both options are favourable for the market?* But if this is our question, the above objection simply does not apply. To put it simply, Sophie already *is* able to “do what she wants”, but her ability to be aware of this fact has been impaired.

3 Teleological Graphs

In “Alienation”, Jaeggi discusses the problem of the so-called “teleological cycle”. Examining this topic provides further evidence for the suitability of our definition, and brings together the various kinds of value we have discussed. It also helpfully highlights some of the most problematic aspects of being epistemically self-alienated. Therefore, I would like to briefly examine how this topic is relevant to us.

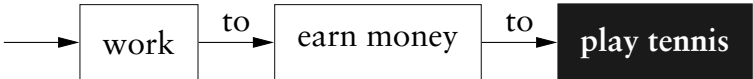
There is a very important difference between objects we value instrumentally on the one

hand, and objects we value intrinsically (or care for) on the other. If we value an object instrumentally, we only value it in relation to some other thing. For example, I find earning money instrumentally valuable because it allows me to play tennis. Therefore, a relation to some other thing is always part of instrumental value — if an object is a means to an end, we must consider the object, but also the end. On the other hand, objects we value intrinsically are defined as being valued in this way precisely by the absence of an external end, and so do not involve an additional relation. The same is true for things we care about — we care about them without regard to some other object.

This gives objects we value instrumentally the additional peculiarity of being able to form chains of reasons: I go to work *in order to* earn money *in order to* play tennis. I will call these kinds of chains of reasons *teleological chains*.


Let us consider different kinds of such chains. For brevity, I will refer to “instrumentally valuable states”, “intrinsically valuable states” and “cared-for states”. By this, I simply mean that one finds oneself in a situation in which one is able to have a positive connection with, or otherwise obtain and enjoy, the object of one’s value or care. If I care about tennis, a cared-for state might be a state in which I can play tennis.

Say I work, in order to earn money, in order to play tennis, because I care about tennis. This is a chain that terminates in a cared-for state. If we draw the instrumentally valuable states as white and the cared-for states as black, we can illustrate this situation as follows:




It is interesting to note that instrumental value is also transitive. If I work to earn money, and earn money to play tennis, it is also reasonable to say that I work to (ultimately) play tennis.

Say we did not care about playing tennis, but still found it intrinsically valuable. This would be a chain ending in an intrinsically valuable state. If we draw intrinsically valuable states with grey, this chain looks as follows:



What if our graph had no intrinsically valuable or cared-for state? In trying to draw such a graph, we need to remember that every instrumentally valuable state must “point to”

something. If there is no non-instrumentally valuable state, we must, at some point, have an instrumentally valuable state point to another instrumentally valuable state, giving us a cycle. This *teleological cycle* means that we do A in order to do B in order to do A.




There are two issues to consider here. Firstly, if we made our decisions according to such a graph, what would our ultimate goal be? We would still have a set of reasons for our decisions, but our planning would not serve anything that we actually care about. I believe that this is what Jaeggi means when she says that “instrumental action (...) is meaningful only if it passes over into an action that leads to a goal that is not itself another means to an end but is pursued for its own sake” (Jaeggi, 2014, p. 208). Therefore, the above arrangement is an example of non-meaningful instrumental action.

But the fact that the above graph leads us to non-meaningful action is not the only problematic aspect. We should also consider that its very cyclical nature gives us the illusory *appearance* that we are doing something for an ultimate purpose. It is obvious that we can tell the difference between a cyclical and non-cyclical graph when considering neat illustrations such as the above. But when we are actually performing an action, we are, so to speak, in one of the states of our graph, and it is easy to only consider the decisions immediately ahead of us. Unfortunately, on a local level, a long instrumental chain with an intrinsically valuable state at the end, and a potentially infinitely long instrumental cycle, look indistinguishable. When faced with such a meaningless cycle, this illusion leads to us expending our energy on projects in hopes of obtaining something that we care about, though no such thing can happen as a result of these particular projects. Jaeggi describes this as follows:


Even in playing the piano, something I do for its own sake, I must do finger exercises that I perform only because of their relation to the external end (external to the exercises themselves) of acquiring the technical skills to play a Beethoven sonata. The possibility of self-realization, however, is threatened precisely when one gets caught in a teleological circle, a situation in which one does one thing only for the sake of another without ever connecting them to a final end, that is, to an end where one can no longer ask the question for what purpose I am doing this? When this characterizes an entire life, the

result is a fatal structure (...) Hence, although instrumental action is always a part of the pursuit of life goals, it is meaningful only if it passes over into an action that leads to a goal that is not itself another means to an end but is pursued for its own sake. (Jaeggi, 2014, p. 208)


So far, we have only considered individual teleological chains relating to specific matters or projects. But what is crucial is to investigate what Jaeggi means when she says that “[w]hen [a teleological cycle] characterizes *an entire life*, the result is a fatal structure” (emphasis mine). Picture a life where one has various projects, each beginning with instrumentally valuable states (light circles) and leading to some cared-for state (dark circles).



The most harmful arrangement is a corruption of the one above. It is when instrumentally valuable states within or across various projects form cycles that define an entire life. This means that, while intrinsically valuable states may still exist, they are simply not reachable from the graph of our reasons, rendering us unable to access them. Since we said that a life where we prioritise things that we care about is desirable, and since we have ended up unable to access any of those things, we can only conclude this is a most undesirable life. Here is an illustration of this, which importantly shows not only the cycles between instrumentally valuable states, but also that the cared-for states (assuming they exist) are not reachable at all:



What would a life that is not at all self-alienated look like? We could illustrate this by drawing a graph with only cared-for states and no instrumentally valuable states:



It is quite straightforward to see that such a life is practically impossible — there is no real way to occupy ourselves with only activities we care about. Daily life is filled with activities that are desirable for us to do but that are not cared for nor even intrinsically valuable. This means that every life is self-alienated to a certain degree, and our hope must be to minimise this self-alienation.

I propose that we can use our definitions of teleological chains to express an alternative (but equivalent) form of our definition of epistemic self-alienation as follows. One leads a non-epistemically-self-alienated life to the extent to which one is epistemically able to construct a teleological graph that allows one to easily access cared-for things. On the other hand, one is epistemically self-alienated to the extent to which one's projects render cared-for things inaccessible. This means that epistemic self-alienation always occurs on a spectrum, since the things we care about can be more or less accessible.

What does it mean for things we care for to be *accessible* or not? Firstly, an object is inaccessible if our decisions cannot lead us to it, which is obvious enough. But it also seems that we are less alienated if we can find ourselves in cared-for states by taking fewer steps, and more alienated if we must follow a very long chain of instrumentally valuable steps before being able to reach a cared-for state. At the same time, this description is not

quite right. These instrumentally valuable steps could be somewhat arbitrarily chosen. The instrumentally valuable act of “working” could be split up into “going to work” and “doing one’s job”, both of which could further be split into their component steps, and so on. Therefore, the number of instrumentally valuable steps is irrelevant.

This issue is compounded by the fact that our teleological graph is not the same as our *realisation* of our teleological graph. For example, I might have a chain consisting of instrumentally valuable A and cared-for B, but this does not mean that B will ever be reached. I might still be “stuck” in A, so to speak, which means I will still be *deprived* of the cared-for B even though it is in my graph. In a sense, my graph is not in itself problematic, but the way that I have realised it is.

I admit I am not sure what the best solution to this problem is, but perhaps it would suffice, for now, to say that cared-for things are accessible to the extent to which we actually spend time in the cared-for states in our teleological graph. That is, the more we actually engage with cared-for things, the better.

Overall, this discussion of teleological cycles strongly supports the definition we have reached. It seems that a paradigmatic “fatal structure” of self-alienation for Jaeggi is the exclusive prioritisation of instrumentally valuable action. The analogy to Sophie’s situation is clear, when we consider that she prioritises instrumentally valuable consultancy work but neglects the sewing that she cares for. This suggests that the antidote to such a situation would be prioritising cared-for things, in order to break out of this teleological cycle. But this is exactly what our definition states: one is epistemically self-alienated if one cannot prioritise cared-for things. Therefore, our definition seems to be compatible with at least this aspect of Jaeggi’s view.

4 Epistemic Blindness

Trigger Warning: This section contains a brief discussion of sexual assault.

I have said that one is epistemically self-alienated to the extent to which one is epistemically blind to the possibility of prioritising things that one cares about. The best description of epistemic blindness comes from Jaeggi herself: “[t]his does not merely mean that he has not acted, or has not availed himself of his possibilities for acting, but that he has not even understood his situation as one in which action is called for or possible” (Jaeggi,

2014, p. 56). This epistemic blindness bears striking similarities to Miranda Fricker's idea of epistemic injustice. If the conception of epistemic blindness underlying our definition of epistemic self-alienation were connected to Fricker's ideas, this might allow us to understand both self-alienation and epistemic injustice better. I would like to argue that epistemic injustice is, in fact, an important component of epistemic self-alienation, and that viewing self-alienation through this lens allows us to better understand the depth of its harmful effects.

Fricker describes epistemic injustice as "a wrong done to someone specifically in their capacity as a knower" (Fricker, 2007, p. 1). What she means by this will be most clearly illustrated by discussing the two kinds of epistemic injustice she describes: testimonial injustice and hermeneutical injustice.

Testimonial injustice "occurs when prejudice causes a hearer to give a deflated level of credibility to a speaker's word" (ibid., p. 1). Fricker's paradigmatic example comes from *The Talented Mr. Ripley* (Minghella et al., 1999), where Greenleaf tells Marge "Marge, there's female intuition, and then there are facts" (ibid., p. 130). Greenleaf exercises a kind of identity power by appealing to a harmful collective social conception that views women as being insufficiently rational and excessively emotional (Fricker, 2007, p. 15). In this situation, Marge's credibility is lessened and her experience is erased, simply because of a class that she belongs to.

On the other hand, hermeneutical injustice "occurs (...) when a gap in collective interpretive resources puts someone at an unfair disadvantage when it comes to making sense of their social experiences" (ibid., p. 1). Fricker refers to Susan Brownmiller's memoir for an example involving a distinguished male professor and a female university employee (Brownmiller, 1990, p. 182). The professor corners the employee and forces her into unwanted physical contact and kissing, leaving the employee traumatised and ultimately causing her to quit her job. When applying for unemployment benefits, the employee is unable to find the words to describe what happened to her, because she does not have a concept that encapsulates her experience, leading to her unemployment claim being denied (Fricker, 2007, p. 149). Fricker describes how one might even struggle to do any better than using clearly insufficient and reductive language to describe such an experience, such as saying it was merely "flirting" (ibid., p. 153). It's only when the employee is able to discuss the incident with others who have had similar experiences that she and

the others are able to put their finger on what happened — *sexual harassment* (Fricker, 2007, p. 150). Hermeneutical injustice therefore causes a “cognitive disablement” that prevents an agent from understanding “a patch of her own experience” that it is “strongly in her interests to understand”, leaving her “deeply troubled, confused, and isolated, not to mention vulnerable to continued harassment” (ibid., p. 151).

It will be useful to note that epistemic injustice in general can take either incidental or systemic forms. If the police does not believe a woman because she is Black, this is an incidental epistemic injustice, because it occurs within a specific interaction. However, such prejudice can often transcend individual interactions. Systemic testimonial injustices “are produced not by prejudice simpliciter, but specifically by those prejudices that ‘track’ the subject through different dimensions of social activity — economic, educational, professional, sexual, legal, political, religious, and so on” (ibid., p. 27). If someone with a Yorkshire accent is viewed as less credible in various spheres of society because of his accent, this is a systemic epistemic injustice, for which there is no one perpetrator.

We have seen that there is clearly an epistemic component to Sophie’s case, and to epistemic self-alienation in general. But can epistemic self-alienation be described as an epistemic *injustice*? Fricker describes how not all hermeneutical disadvantages are necessarily injustices:

Different groups can be hermeneutically disadvantaged for all sorts of reasons, as the changing social world frequently generates new sorts of experience of which our understanding may dawn only gradually; but only some of these cognitive disadvantages will strike one as unjust. For something to be an injustice, it must be harmful but also wrongful, whether because discriminatory or because otherwise unfair. (ibid., p. 151)

It is clear, and helpful, to note that some hermeneutical disadvantages are problematic, and some not. However, it is inconvenient to base our criteria for why epistemic self-alienation is epistemically problematic on Fricker’s description because, if we did so, we would have to disentangle various notions Fricker uses. She refers to *harm*, *wrongfulness*, *injustice* and also *unfairness*, but it is not immediately obvious how these concepts should all be defined, and how they interact and overlap with each other.

Therefore, for my part, I will simply say that epistemic self-alienation is sometimes epis-

temically *wrongful*. Sophie is wronged because social coercion causes her to have a defective epistemic relationship to something that she should have a positive epistemic relationship to, namely the process of planning her life journey. I will continue referring to epistemic injustice as Fricker describes it, but when I say that epistemic self-alienation is an epistemic injustice, I ultimately mean that this epistemic wrongfulness is present.

Can we, then, still say that this wrongfulness is compatible with what Fricker means when she refers to epistemic injustice? I believe so. Fricker describes exactly the phenomenon of social coercion I have referred to as one in which “prejudicial stereotypes in the social atmosphere” (Fricker, 2007, p. 153–154) harm an agent, either because of the agent’s identity or because of their socio-economic background. This seems like a good match for our example, because Sophie is exposed to exactly this kind of prejudice — her desired career as a garment maker is viewed as “not a real job”, discouraging her from exploring this possibility for fear of further epistemic marginalisation.

In this way, Sophie is wronged in her capacity as a knower. It seems clear that this is an “injustice” according to Fricker’s conception, and the aspect of hermeneutical marginalisation certainly suggests that Sophie is experiencing a hermeneutical injustice. Even more encouraging is Fricker’s description of what happens when one overcomes hermeneutical marginalisation. She says that “women were collectively able to overcome extant *routine* social interpretive habits and arrive at *exceptional* interpretations of some of their formerly occluded experiences” (ibid., p. 148), which is obviously analogous to our previous discussion of epistemic self-alienation in terms of “unique outcomes” as opposed to the “dominant story”.

While this is very encouraging, there is an argument we should consider for why Sophie’s experience might *not* be a hermeneutical injustice. Fricker’s instances of hermeneutical injustice seem to involve someone *lacking* a certain hermeneutical resource, or to put things more simply, a certain concept. In the sexual harassment example, the victims lacked the concept of sexual harassment, which impaired them from making sense of their experience. It does not seem that Sophie’s situation is caused by her lacking a concept. For example, both career paths are conceptually open to her: she has the concept of *consulting-company-career*, as well as the concept of *sewing-career*. Are we, then, really dealing with a hermeneutical injustice?

This objection is very helpful, because it encourages us to clarify an important part of

Fricker's definition, and this clarification is, in fact, my main point. Remember that Fricker says that hermeneutical injustice "occurs (...) when a *gap in collective interpretive resources* puts someone at an unfair disadvantage when it comes to making sense of their social experiences" (emphasis mine, Fricker, 2007, p. 1). What, then, is "a gap in collective interpretive resources"? One possibility for such a gap is the lack of a certain concept, but are there other possibilities?

My claim is that there exist other kinds of such gaps. Particularly, it is possible to be acquainted with a concept while simultaneously being blind to the possibility of relating to, or apprehending, that concept in a certain way. It may be that Sophie is familiar with the concept of *one* having a career at a consulting company and that of *one* having a career in sewing, while nevertheless not being able to apprehend the possibility of *her* having a career in sewing. Therefore, what is made inaccessible by structural prejudice is not the concept of a sewing career, but the possibility of Sophie engaging with a sewing career in a certain way. My suggestion is to bring the following addendum to Fricker's definition of hermeneutical injustice:

Definition 2 *A gap in collective interpretive resources can occur not only when one does not have a concept, but also when one does have a concept, but cannot conceive of it being applicable to oneself or one's own life.*

There is another potential objection, which leads us to another helpful point. Fricker seems to believe that hermeneutical injustice chiefly happens *communicatively*, in relation to another. The main point in the sexual harassment example was that the women could not convey to others that they had been sexually harassed, because this shared hermeneutical resource did not exist: "the primary harm of hermeneutical injustice consists in a situated hermeneutical inequality: (...) the subject is rendered unable to *make communicatively intelligible* something which it is particularly in his or her interests to be able to render intelligible" (emphasis mine, *ibid.*, p. 162). But in our example, Sophie's main difficulty is not that she cannot communicate something to someone else, but that she cannot understand herself and her possibilities of choosing. Her difficulty is not *communicative*, but *reflective*. Might this, then, be a reason to say that Sophie is not experiencing hermeneutical injustice?

In fact, the notion that hermeneutical injustice occurs only communicatively and not reflectively is mistaken, and a more detailed reading of Fricker's comments shows this:

The primary harm of hermeneutical injustice, then, is to be understood not only in terms of the subject's being unfairly disadvantaged by some collective hermeneutical lacuna, but also in terms of the very construction (constitutive and/or causal) of selfhood. In certain social contexts, hermeneutical injustice can mean that someone is socially constituted as, and perhaps even caused to be, something they are not, and which it is against their interests to be seen to be. Thus, as we put the point previously in our discussion of the wrong of testimonial justice, *they may be prevented from becoming who they are.* (emphasis mine, Fricker, 2007, p. 168)

If an impaired construction of selfhood can also be the primary harm of hermeneutical injustice, and a construction of selfhood is not necessarily inherently communicative in nature, it stands to reason that it is possible for hermeneutical injustice to operate chiefly reflectively.

Definition 3 *Hermeneutical injustice can occur both communicatively, when an agent is impaired in her communication with other agents, and reflectively, when an agent's understanding of herself is impaired.*

What, then, is the relationship between epistemic self-alienation and epistemic injustice? Recall that we defined self-alienation as follows:

Definition 1 *One is epistemically self-alienated to the extent to which one is epistemically blind to the possibility of prioritising things that one cares about.*

I claim that the epistemic blindness we have been referring to is itself an instance of hermeneutical injustice. Our definitions seem to align in suggesting this. Epistemic blindness can occur because of a gap in collective interpretive resources. This gap can be present when the concepts themselves are known to the agent, but an agent's relationships to these concepts are impaired — we refer not to the agent being blind to the things that one cares about, but to the possibility of prioritising them, i.e. relating to them in a certain way. And, as we have seen, there is no difficulty in the fact that this epistemic blindness occurs reflectively and not communicatively. Importantly, both the epistemic blindness of epistemic self-alienation and the symptoms of hermeneutical injustice result in an impaired conception of self. It is notable that both Jaeggi and Fricker describe their respective concepts of self-alienation and hermeneutical injustice in a similar way: they

stop one from being who they are (Jaeggi, 2014, p. 99; Fricker, 2007, p. 168).

What then, is the relationship between epistemic self-alienation and epistemic injustice? I claim that the epistemic blindness that is constitutive of epistemic self-alienation is itself an epistemic injustice. Therefore, epistemic injustice is a symptom of epistemic self-alienation.

Definition 4 *Epistemic self-alienation causes hermeneutical injustice.*

We have seen that hermeneutical injustice can be incidental or systemic. Is the hermeneutical injustice caused by epistemic self-alienation incidental or systemic, or can it be both? I believe that it is always systemic, simply because epistemic self-alienation does not generally occur as the result of a singular interaction, but rather because of a web of socially inherited prejudices and ideas.

It is fully reasonable to ask what, if anything, we have gained by reaching this definition. Have we simply shuffled around concepts with little benefit? To be sure, my discussion has mostly dealt with definitions, but I believe that highlighting such a powerful relationship between self-alienation and epistemic injustice can not only allow our understanding of one to inform that of the other, but also to highlight the harms of self-alienation. While I believe that many useful analyses can be performed with the aid of this connection, for now I will simply provide an illustrative example of such a harm.

Hermeneutical injustice causes different kinds of practical harm. In Sophie's example, her career, and therefore her life's path, is constrained by the prejudice imposed upon her. But hermeneutical injustice can also cause a very different kind of harm, namely what Fricker calls a secondary epistemic harm. The clearest example is a loss of epistemic confidence. The women who were sexually assaulted in the above example knew deep down that something wrong had happened, but the dominant story did not allow them to understand it, and this dissonance between one's "intimated sense of a given experience" (ibid., p. 163) and the dominant story can erode one's faith in one's own epistemic capabilities. Sophie knows that a career in sewing would be meaningful to her, but the dominant story is that it is "not a real job". In this situation, Sophie might conclude that she is not good at reasoning about career options, causing her epistemic confidence to be reduced.

This loss of epistemic confidence can cause serious harm. Agents who are not convinced of their capacity as knowers can either lose existing knowledge, or be impaired in the acquisition of new knowledge. Worst of all, this can result in a negative feedback loop and a self-fulfilling prophecy — the more an agent’s epistemic confidence is impaired, the more they will fail to have a good relationship to new or existing knowledge, thereby only confirming their prejudice about themselves, causing further loss of knowledge, and so on.

The harm caused here is clearly significant, but what is notable here is that, as a result of our findings, we can say that this harm is not only caused by epistemic injustice, but also by (epistemic) self-alienation. Our definitions have therefore shone a light on the possible epistemic nature of self-alienation, and also on the extent of the epistemic harm it causes.

5 Conclusion

In this paper, I have attempted to understand Sophie’s example better by introducing and defining the idea of epistemic self-alienation, and inspecting this condition through an analysis of the concepts of self-alienation and epistemic injustice. In this endeavour I have reached the following four conclusions:

Definition 1 *One is epistemically self-alienated to the extent to which one is epistemically blind to the possibility of prioritising things that one cares about.*

Definition 2 *A gap in collective interpretive resources can occur not only when one does not have a concept, but also when one does have a concept, but cannot conceive of it being applicable to oneself or one’s own life.*

Definition 3 *Hermeneutical injustice can occur both communicatively, when an agent is impaired in her communication with other agents, and reflectively, when an agent’s understanding of herself is impaired.*

Definition 4 *Epistemic self-alienation causes hermeneutical injustice.*

I hope that the problematic aspects of Sophie’s situation have been made clearer through my analysis. I believe this is important because Sophie’s example is not an arbitrarily constructed one, but a very real one, that touches and affects almost everyone in Western

society ⁷ in one way or another at some point in our lives. Sophie's situation is one we have all encountered in ourselves and in others, and yet it is one that causes serious harm. I think that even a slight contribution towards a better understanding of this situation is important.

Opportunities for further work are plentiful, and one stands out in particular. We have seen that the hermeneutical injustice caused by epistemic self-alienation is systemic. There are certainly aspects of our social institutions, such as families, local communities and corporations that perpetuate and enable such prejudices. Investigating what these aspects are might provide fruitful opportunities to mitigate these injustices. ⁸

References

- Adorno, Theodor W (1973). "Negative dialectics". Vol. 1. A&C Black.
- Brownmiller, Susan (1990). "In Our Time: Memoir of a revolution". New York: Dial Press.
- Foucault, Michel and Elissa Mailänder (1975). "Surveiller et punir". Vol. 225. na.
- Frankfurt, Harry (1999). "Necessity, Volition and Love". In: Cambridge University Press.
- Fricker, Miranda (2007). "Epistemic injustice: Power and the ethics of knowing". Oxford University Press.
- Geertz, Clifford (1986). "Making experiences, authoring selves". In: *The anthropology of experience*, pp. 373–380.
- Hussain, Waheed (2023). "Living with the Invisible Hand: Markets, Corporations, and Human Freedom". Oxford University Press.
- Jaeggi, Rahel (2014). "Alienation". Columbia University Press.
- Kambartel, Friedrich (1989). "Universalität als Lebensform". In: *Philosophie der humanen Welt*.
- Leopold, David (2022). "Alienation". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Uri Zalta Edward N.; Nodelman. Winter 2022. Metaphysics Research Lab, Stanford University.

⁷These problems are certainly present in other societies, but I have limited my claim to Western society so as to avoid making claims about cultures I am unfamiliar with.

⁸I am grateful to Han van Wietmarschen for his supervision and for his helpful and thorough feedback on this paper. I am also grateful to the following for their feedback and advice: Allyson Obber, Daisy O'Connor, Joseph Al-Khalili, Tian Lan, Tom Stern. I thank the following for their unwavering support: Harriet Boeing, Harry Hunt, Liam Barer.

- Lewis, David (1983). "Extrinsic properties". In: *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 44.2, pp. 197–200.
- Lukács, Georg (1971). "Reification and the Consciousness of the Proletariat". In: *History and class consciousness: Studies in Marxist dialectics* 137, pp. 83–222.
- Marcuse, Herbert (2013). "One-dimensional man: Studies in the ideology of advanced industrial society". Routledge.
- Marx, Karl and Friedrich Engels (1965). *The German Ideology (1845)*.
- Minghella, Anthony et al. (1999). "The Talented Mr. Ripley". Paramount Pictures Hollywood, CA.
- Stoljar, Natalie (2000). "Autonomy and the Feminist Intuition". In: *Relational autonomy: Feminist perspectives on autonomy, agency, and the social self*. Ed. by Natalie Mackenzie Catriona; Stoljar. Oxford University Press.
- Weber, Max (1978). "Economy and society: An outline of interpretive sociology". Vol. 2. University of California Press.
- White, Michael and David Epston (1990). "Narrative means to therapeutic ends". WW Norton & Company.
- Wildt, Andreas (1987). "Die Anthropologie des jungen Marx".