



Universiteit
Leiden
The Netherlands

On Responsibility for Structural Injustice

Bramer, Mickey

Citation

Bramer, M. (2023). *On Responsibility for Structural Injustice*.

Version: Not Applicable (or Unknown)

License: [License to inclusion and publication of a Bachelor or Master thesis in the Leiden University Student Repository](#)

Downloaded from: <https://hdl.handle.net/1887/3589733>

Note: To cite this publication please use the final published version (if applicable).

On Responsibility for Structural Injustice

MA Thesis in Moral and Political Philosophy

Leiden University, Institute for Philosophy

Supervised by Dr. Wouter Kalf

Mickey Bramer

January 15th, 2023

Table of contents

Table of contents	2
Introduction	3
Chapter One: Conceptualizing structural injustice	6
1.1 Social structure	6
1.2 Structural injustice	8
1.2.1 <i>Three preliminaries</i>	9
1.2.2 <i>The causal claim: structural factors as a causal explanation</i>	10
1.2.3 <i>The evaluative claim: wrongness</i>	11
1.3 Conclusion.....	12
Chapter Two: Conceptualizing responsibility for structural injustice	13
2.1 Two different ideas of responsibility, divided into four types.....	14
2.2 Evaluating the responsibility types in the context of structural injustice	16
2.3 Young’s distinction: The Liability Model and the Social Connection Model	17
2.4 Counterarguments extending moral responsibility to structural injustice.....	18
2.4.1 <i>The Harmless Torturers</i>	18
2.4.2 <i>Complicity</i>	19
2.5 Incoherence in time-dimension.....	20
2.5.1 <i>The challenge</i>	20
2.5.2 <i>Two new questions</i>	21
2.6 Conclusion.....	22
Chapter Three: Allocating remedial responsibility	24
3.1 Reproducing structural injustice as relevant ground for remedial responsibility	24
3.2 Arguments from capacity to reject participation.....	26
3.3 Overdemandingness.....	27
3.4 Conclusion.....	28
Conclusion	29
Acknowledgements	29
Bibliography	30

Introduction

In today's globalized world, many problems we face are large-scale and complex. Some examples are: global as well as national poverty; the current climate crisis, which will result in massive biodiversity loss and will cost millions of lives in the near future according to the latest IPCC rapport¹; inequality and oppression of minority groups. I am not interested in these specific problems, but in the general *genus* of problems of which I believe they are *species*: they tend to have complex causalities, are multi-dimensional, involve large numbers of actors with different interests, and often cross over administrative and generational boundaries. As such they raise pressing political-philosophical questions.

In this thesis I want to investigate this category of problems from the perspective of responsibility, with a particular focus on individual actors. To quote Iris Marion Young, whom I will discuss at length later: "How should we as individuals think about our own responsibility in relation to [these problems]?" (2011, p.15)² Everyday life shows us how people have different ideas about this. On the one hand, for example, we see civil mass action movements like Extinction Rebellion and Greenpeace, who believe individuals have a responsibility regarding the climate crisis. Another often-cited example in this context is the anti-sweatshop movement, building on the intuition that consumers in western societies (who buy clothing produced in sweatshops) are somehow responsible for the working conditions in those sweatshops. (Young, p.125-134; Barry & Macdonald, 2016, p.92) But on the other hand, there are arguments to the contrary, claiming that responsibility lies elsewhere, e.g. with polluting corporations or at the bosses of sweatshops who actively create the problems. Individual citizens merely mind their own business and cannot be held responsible for all the big problems in the world.

I want to know which claims on this matter (if any) are correct. I will approach this inquiry in the following three steps. Firstly, I conceptualize the genus of the problems of interest as *structural injustices*: injustices that are causally explained by reference to social structure. I am interested in responsibility for this specific type of injustices, and so we first need to develop an understanding of their characteristics. Then, secondly, we need to ask: what does the word 'responsibility' *mean* in this context? I conduct a conceptual analysis of the term, where I differentiate between different types of responsibility and argue that the structural context forces us to focus on one of these. Then thirdly, I investigate the *grounds* of this responsibility type. In other words, what are the properties or conditions that generate such a responsibility? Determining *who* is responsible is then only a matter of identifying which agents satisfy those conditions.

To clearly demarcate the line I set out here, it is useful to consider the questions I do *not* focus on, as the topic invites many related questions. I touch on some of them throughout the thesis but only briefly, to the extent that is required for the steps above. To start, allocating responsibility is only done in the abstract; the conclusion will be of the type 'all agents who have property X', leaving open who these agents are specifically. Moreover, I will only point in a very broad direction on the *content* of the responsibility (i.e. what conduct it prescribes), as opposed to concrete, specific actions. This is due to space constraints. The question of grounds has priority over that of content: it seems more useful to me to know who is responsible and only broadly what must be done, than to know what to do exactly but not who must do it. I also do not address the practice of *holding* others responsible in this context, i.e. what we can do if we see others fail to act on their responsibility. Again, that becomes more appropriate in a later stage, once we know who is responsible for which actions.

Having established the general outline, why is this project relevant philosophically? Firstly, in the words of David Miller, "few concepts [...] are more slippery" than responsibility. (2001, p.455) This makes a conceptual analysis interesting, challenging, and potentially rewarding as a contribution to the field. Secondly, as I am interested in structural injustice as a category of problems, any conclusions reached here will, *mutatis mutandis*, also be applicable to specific elements of the category. Some alleged structural injustices have lately attracted substantial philosophical attention due to their

¹ <https://www.ipcc.ch/report/ar6/wg2/>, accessed on 13-01-2023.

² All references to Young are from *Responsibility for Justice* (2011). Henceforth I will omit the year when citing her in-text.

urgency, such as the climate crisis leading to a surge in environmental philosophy. It is hoped that this thesis might also be of value within those more specific research areas. Thirdly, I am not the first to engage with these questions. The topic rose to prominence through Young's seminal work *Responsibility for Justice* (2011), which attracted many commentators over the last decade, among others Maeve McKeown (2015, 2018, 2021), Robin Zheng (2018, 2019, 2021), Andrea Sangiovanni (2018) and Mattias Gunnemyr (2020). They have shown how Young's theory provides valuable insights while also failing in some regards. However, most commentators only address one element of Young's total account (which is comprehensive): they each highlight a defect and try to solve it with their specific contribution, like changing a part of a half-working car. But as this number of commentators has grown, so too has the number of spare parts lying around. My hope here is to assemble these parts into an improved whole, while still honoring the spirit of Young's original work.

What exactly will I argue? I claim that, somewhat like Young, *participation in social structures that cause injustice generates a pro tanto responsibility to combat that injustice*.³ Let me clarify this claim, as well as how it relates to the existing literature, by looking at the individual chapters.

In the first chapter I break with most of the literature and *explicitly* discuss the legitimacy of 'structural injustice' as an analytical tool: it allows us to identify previously unnoticed injustices or describe already noticed injustices in a more accurate way. I make these steps explicit because the other chapters build on this: we can only convincingly discuss responsibility for structural injustice in full knowledge of the latter. The beginning of the chapter discusses the workings of social structure. Social structures constrain the actions of the people within it, while simultaneously only existing thanks to large masses of people acting in accordance with them: they are recursive upon our actions. I then use a broadly accepted theory of causality to show how structures can cause injustice. For clarity I contrast this with 'interactional injustice' (a term I borrow from Catherine Lu, 2017), which is what we typically think of when we think of injustice: a clear perpetrator inflicting harm on a victim. All this spells out the following conclusion: millions of agents act and reproduce structures, and those structures cause injustice. It is those millions of reproducing actions that, as I argue in the third chapter, generate responsibility.

The second chapter then offers a conceptual analysis of responsibility with respect to injustice. I depart from Young and follow David Miller's analysis, which I defend as being more accurate. He identifies two fundamental but distinct ideas of responsibility, under which a total four types of responsibility fall. Idea 1 concerns the link between a person's agency and an outcome, involving concepts like fault, guilt and blame. Under this fall causal, outcome and moral responsibility. Idea 2 concerns the burden of correcting a bad situation: remedial responsibility. While Idea 1 dominates discussion in dealing with interactional injustice, it seems largely inapt in dealing with structural injustice: I argue that only partial causal responsibility can be established in the structural context, and not outcome or moral responsibility. Instead, I argue that the structural context forces us to focus on remedial responsibility. But this does not yet tell us anything about which conditions or properties generate such a responsibility - that will be examined in the third chapter.

But before we can turn to that, the remainder of the second chapter deals with two challenges. The first consists of arguments that claim we *can* use Idea 1 in the structural context. I discuss and reject two of these arguments. The second challenge is internal to Miller's distinction (and has also been voiced against Young): it wonders whether it is at all coherent to speak of a responsibility type that is forward-looking *only* (thus without terms like blame and guilt). Because what then if someone fails to take responsibility? I give a different response than Young and argue that this objection is based on a misunderstanding.

Finally, then, the third chapter establishes the conditions that determine who has remedial responsibility for structural injustice. I claim that the millions of actions that reproduce the harmful structures (as discussed in chapter one, and identified as partial causal responsibility in chapter two) generate remedial responsibility in the structural context. Even though an individual agent is not guilty for the structures, because one's influence over their existence is highly limited, participating in

³ I use the word 'somewhat' because Young has not unequivocally embraced this conclusion: she speaks of 'social connection' to the injustice that generates responsibility, which has mostly (and rightly, I believe) been interpreted by others as intending participation (see McKeown, 2018 and Gunnemyr, 2018).

them still has moral weight. What we do is this: even though we focus on Idea 2 of responsibility, we take a minor element of Idea 1 and use this as a *criterion* for allocating Idea 2.

The remainder of the chapter considers another possible contender for the responsibility generating criterion: the capacity to solve the injustice, often illustrated with the famous thought experiment of the drowning child in a pond. I concede that capacity could indeed ground remedial responsibility, but such a capacity almost never exists in the structural context. Even if it does, capacity generates responsibility by virtue of some general humanitarian principle. This is not specific to the structural context. Any responsibility generated in this way is simply another (*pro tanto*) responsibility that complements the structural-specific responsibility generated by participation. The upshot of the discussion on capacity is that there exist circumstances in which remedial responsibility for structural injustice dissipates: it can be outweighed by other responsibilities (as it is *pro tanto* as opposed to *all things considered*) or even excused altogether.

In sum: the aim of this thesis is to establish participation in harmful social structures as the ground for remedial responsibility for structural injustices. I claim that the reproduction of structures does not warrant blame or guilt, but bears some moral weight nonetheless.

I want to conclude this introduction by resisting a preliminary worry: why do we have to do all this philosophical gymnastics with responsibility, and can we not simply allocate remedial responsibility on pragmatic grounds? We could set certain goals of (social) justice and then wonder to whom we could best assign responsibility to achieve such goals. I reply that we do not just want to solve the structural injustices claspings our world, we want to solve them *fairly*. Miller (2007) distinguishes between *identifying* responsibility, which is a search for morally relevant grounds (like my project here), and *assigning* it, which is realizing a certain responsibility distribution out there in the world. I take it that any fair responsibility distribution must be at least somewhat based on the identified responsibility; we want to hold the right parties responsible for the right reasons.

Chapter One

Conceptualizing structural injustice

This chapter is concerned with conceptualizing the problems that this thesis is interested in as “structural injustices” (called ‘SI’ in the bulk of the text). The term was popularized by Young, leading to Robin Zheng even speaking of a “turn to structure” in recent thinking. (2018, p.271) The concept mainly serves as a useful analytical tool (Sankaran, 2021): it opens up possibilities for recognizing and describing injustices where our traditional toolbox falls short. Examples of topics that have already been analyzed structurally are, among others, international justice and colonialism (Lu, 2017), the climate crisis (Eckersley, 2016), immigration (Parekh, 2012) and health vulnerability (Chung, 2021). Now for this thesis I am interested in (responsibility for) SI as a general *genus* of problems, of which I take these examples to be more specific species. To say that a problem is an SI is to say that it is (at least partially) causally explained by socio-structural effects. This chapter establishes how such causation works.

I make this structural causation explicit for multiple reasons. Firstly, the plausibility of a ‘turn to structure’ is not at all obvious, as structural explanations “resist an easy grasp” in comparison to individualistic explanations. (Ayala-Lopez, 2018, p.727) The ontological status of the main element remains vague (the individual versus social structure; ‘what is a structure?’), and the same goes for its explanatory role. These questions should be addressed before a discussion on responsibility in the structural context is even warranted. Secondly, in the next chapter we are not just looking for any conception of responsibility, but one that is *specifically* attuned to the structural context. We can only formulate this convincingly in full knowledge of the latter. Thirdly, leaving these matters implicit (like many authors do) potentially obscures the debate. It seems to me that sometimes authors disagree with each other on their ideas of responsibility because they tacitly seem to have a different understanding of SI’s and their causal nature. In sections 2.4.1 and 3.2 I consider what I think are examples of this. Spelling out the characteristics of SI’s here helps us resolve arguments later on.

Finally, let me give a brief overview of the structure of this chapter. In the first section I discuss what social structures are and how they function. In Sangiovanni’s work (2018) I find a general and abstract picture of social structure, whereas Young’s treatment is opposite, providing specific qualities but no general framework. To fully appreciate the workings of structures I believe we need both.⁴ In the second section I connect social structure to injustice. I use a broadly accepted theory of causality to show that - if that theory of causality is correct - one must indeed appeal to structures in order to explain certain problems. Next to this, to say that something is a structural *injustice* is also, more obviously, to say that it is not merely bad but *wrong*. Something is bad when it is harmful (‘drinking petrol is bad for you’), but a harm is only wrong when it is unfair (‘it was wrong of him to make you drink petrol’). Depending on the circumstances a bad action can be the right one, and vice versa. I briefly discuss this evaluative part of the concept by contrasting SI with mere bad luck. The third section concludes.

1.1 Social structure

Sangiovanni discusses the core mechanism of structures. He invites us to first consider a structure in the abstract, which is a “representation of the set of patterned relations possible for a given set of objects.” (p.461) In the example of a chessboard, this means a list of all the pieces, the board they move on and which moves are allowed for each piece. If we apply this concept to a social environment we get an abstract *social* structure (for example ‘courtly life in the Renaissance’): all the different *social* roles available (the pieces), the background to which they act (the board), and the moves that are

⁴ While I have only used Young and her commentators as sources for this chapter, most of *their* work can be traced back to the sociologists William Sewell, Anthony Giddens and Pierre Bourdieu. I have decided not to cite them directly, because that would shift the emphasis of this thesis too much (as I am not a sociologist), and because I think their interpretation by Young and others is already concrete enough for my purposes.

allowed for each role. A structure in this social sense is more extensive than chess, because which roles and moves are available is tied to an entire cultural apparatus: “rules, concepts, beliefs, assumption, conventions, recipes, scenarios, principles of action, and habits of speech,” which can be both implicit and explicit. (id.) These cultural elements are called ‘schemas’.

These abstract schemas become ‘practices’ when they are enacted in the real world – ‘social structure’ is the entirety of schemas, practices and their interplay. The difference between schema and practice is analogous to the difference between chess and a particular *game* of chess. Which specific form the practice adopts (which particular game sees the board) depends on whether agents have the means to act out their allowed moves: it depends on the distribution of *resources* (roughly defined as “any material or nonmaterial thing that is useful for getting what we want”, p.462). What counts as a resource is itself determined by the schemas, as those define what is valuable, such as ‘being king’ or ‘having oil’. But it works the other way around as well: a certain resource distribution or practice can influence schemas by causing changes in conventions or beliefs. Taking all this together, “social structures [...] are schemas, as enacted in practices by particular people in particular places with particular things, that both are shaped by and give shape to distributions of material and nonmaterial resources.” (id.)

For this very general understanding to be useful for the conceptual analysis of responsibility in the next chapter, we must distill from it some concrete characteristics. The ones that I present here are inspired by but not limited to Young’s. (p.53-64) Firstly, social structures are *constraining* because they enable and disable, favor and disfavor certain modes of behavior. Abstract schemas can do this directly by prescribing certain conduct for certain roles, but also indirectly by shaping the resource distribution. Again, such constraints can be material (e.g. physical, economic) and immaterial (e.g. socially, institutionally). An insightful analogy is infrastructure and traffic: we can only move where roads and trainlines have been built, where the traffic rules allow us to go, and all depending on which means of transportation we can afford in the first place. Notice how ‘constraining’ in this sense is not coercing but steering: we are not forced into one choice, but rather structures “produce differentials in the kinds and range of options that individuals have for their choices.” (Young, p.55)

Secondly, social structures are *recursive* and often *self-reinforcing*. Recursion happens when something is defined in terms of itself (such as the visual Droste effect). This is indeed the case for structures: resource distributions are shaped by schemas, but those schemas are influenced by prior resource distributions, and those are influenced by prior schemas, etc. Returning to the analogy of infrastructure, such recursion is often self-reinforcing: after the road is built people and businesses settle next to it, which makes it increasingly arduous to relocate the road. It reinforces its own existence. Moreover, the more this happens the more arduous it becomes to even *want* to relocate the road. This is the third point: the constraints of social structures often appear as *objective* and tend to evade scrutiny. Both materially (roads) and immaterially (social norms), people tend to experience structures as objectively given. In terms of direct experience, most of us do not consciously consider street plans or social norms, unless we feel bothered by them. They are the background conditions, the fixed point on which we have to build our daily conduct.⁵

Fourthly, social structure only exists by virtue of *large numbers of actions*; it “refers to the accumulated outcomes of the actions of the masses of individuals enacting their own projects, often uncoordinated with many others.” (Young, p.62) They result from actions because actions form the bridge between schemas and practices; it has to be large numbers because no individual agent can shape schemas or practices alone. There is a duality here: even though structures are constraining and seen as objective, they are dynamic and depend on the agency of masses of people. Young tries to emphasize this by referring to *social-structural processes* instead of the more static formulation of ‘structure’. She notes: “when individuals act, they are doing two things at once: (1) They are trying to bring about a state of affairs that they intend, and (2) they are reproducing the structural properties, the positional relations of rules and resources, on which they draw for these actions.” (p.60)

⁵ I see interesting parallels between the concept of social structure here, and continental ideas elsewhere. I think of Marx’s ideas of ideology and false consciousness, describing how the personal world experience is tied to an entire (cultural) apparatus of power relations. I also think of Heidegger’s *Die Frage nach der Technik* (1954), in which technology, more than being of instrumental use, creates an ‘Umwelt’, the surrounding world which constitutes and influences human activity. But of course these parallels are way beyond the scope of this thesis.

Fifthly, the emergence of structures is to a high degree unintentional. Participating agents are simply minding their own business, pursuing their own goals. Even though the structures that result from this may be harmful, most of the actions that align with structural norms are even seen as good, precisely *because* they are aligned. Many people would even regret the existence of some harmful structure they participate in, were they asked about it. This unintentionality becomes more convincing once we consider the complexity of most structures: they often involve economic, political, cultural and physical schemas and resources (think again of the Renaissance courtly life). Intentional creation of structure would mean a superhuman level of ability and commitment: agents must see their actions in correspondence with the large number of other people doing the same, and correctly evaluate the consequences of this across multiple domains.

To conclude: social structures are constraining in an option-shaping way; recursive and often self-reinforcing; experienced objectively; dynamic, in that they exist in the actions of masses of individuals; and mostly unintentional. These characteristics form the structural context on which the rest of this thesis builds, and will all play their part in the arguments to come.

1.2 Structural injustice

Let us now connect this picture of social structure with the concept of structural injustice. Sangiovanni has pointed out that the term is ambiguous, so we should start by clarifying its meaning. (p.462-3) SI does not say anything about what the injustice consists of, but about what best *explains* the injustice. It is an explanatory concept rather than a substantive one. The latter would say something about which things can be said to have the property of justness or unjustness. SI, by contrast, is “injustice that is explained by the presence of a structure rather than by an aggregate of individual decisions operating independently or in spite of a structure.” (id.) In short, SI’s are injustices caused by structures.⁶ For clarity we may contrast this with the alternative, as formulated in the second half of the quote: injustices caused by actions of individual agents. Let us call this ‘interactional injustice’: the victim is harmed in a clear interaction with the perpetrator.⁷ (Lu, 2017)

To say that something is an SI is also, obviously, to say that it is an injustice, i.e. that the problematic situation at hand is not merely bad but *wrong*. This is an evaluatively loaded claim, as opposed to the causal claim which is evaluatively neutral. To be sure, often causal and evaluative claims go hand in hand – ‘Beth has been murdered by Alex’ arguably leads to a different evaluative conclusion than ‘Beth has been struck by lightning’ – but they may come apart. Thus, in order to argue that a problem is an SI, we must argue for two claims: that social structures play a role in its cause, and that the problem is unjust.

The aim of this section is to elaborate these two claims for one instance of SI, with emphasis on the causal claim. We have to understand the causal nature of SI to adequately conceptualize responsibility for it in the next chapter. But besides this, I want to convince the reader that applying the concept of SI is valid, i.e. that some problems really should be conceptualized as caused by structures. Admittedly, my argument for this is inductive, as I only have space to cover one instance in depth. But given the number of structural analyses already available (see the introduction of this chapter for examples, and examples in the chapters to come), I think it is uncontroversial to assume that my analysis here generalizes to other instances as well. I claim that there exist at least *some* instances of SI; it is not an empty category.

Before considering the two claims, I first discuss some preliminaries, among which a commonly accepted theory on causality by Hart and Honoré.

⁶ Explanation and causation are not necessarily the same. Without going too deep into the metaphysics of explanations, I follow Haslanger (2016) in thinking of explanations as answers to questions. In this case, it is an answer to the question ‘How did this injustice arise?’, to which the answer is a causal statement (about social structures). Sangiovanni glosses over this.

⁷ This can be a person or a collective agent, like a government. In case of the latter we would speak of ‘institutional injustice’, but this belongs to the larger category of interactional injustice. (Aragon & Jaggar, 2018, p.442)

1.2.1 Three preliminaries

First I will illustrate the instance of SI we will be examining in the next section: Young's classic example of the nearly homeless Sandy. (p.43-52) Single mother of two children and without any college degree, Sandy is about to be evicted from her central-city apartment by the landlord, because the building has been generating a loss and he wants to convert it into condominiums. She fanatically searches for a new place, with two desiderata: a reasonable public transport connection to her work (she works in a suburban mall and has no car), and a nice neighborhood for her children with a good school nearby. But her search, so goes the example, fails spectacularly because of multiple factors. The area close to her work is mainly filled with small houses for families, too large and expensive for her; the few apartments that are near are too expensive as well; any apartment she can actually afford brings with it a total commute of four hours a day due to its location; she cannot find a new job in the short time she has before eviction; and any social welfare program she would qualify for has a two-year waiting list. As a final attempt she buys a car with all her savings so that she can drop her first desideratum, only to find out that she now lacks the money for a down payment consisting of multiple months' rent, which almost all landlords require. All these factors combined result in Sandy being on the brink of homelessness.

Before we can establish the causal role of social structure in Sandy's story, we must first consider the concept of causation itself. What does it mean to say that something (social structure, in this case) is a causal explanation of something else (Sandy's position)? This is a much more complicated and extensive question than this thesis can cover; I wish to steer clear of this Pandora's box, while still covering the topic adequately for our purposes. Therefore I will resort to one theory of causality (by Hart and Honoré, 1959), which I take to be uncontroversial: it has been historically influential, relatively accessible in its outset, and intuitively powerful. It has also been used by proponents and opponents of Young alike (e.g. McKeown (2018) and Gunnemyr (2020) respectively).

The question of causality is fundamentally concerned with finding the appropriate relation between some event happening and a certain outcome occurring. Imagine I drop a lit match on a pile of kindle wood, and a fire starts. What is it that makes us say that my match dropping *causes* the fire? A first guess could be that my match was a necessary condition for the fire occurring. This is the basic conception of a cause-in-fact in law (a cause in an explanatory sense), also called the '*sine qua non*'-test: "but for the defendant's action, would the victim have been harmed as she was?" (Moore, 2019) But upon further inspection this test cannot be adequate alone, among other things because it is overinclusive. Not only is my match dropping a necessary condition, but also the oxygen in the air, the pile of kindle wood, etc. This is an unsatisfactory result, because we would generally not consider 'oxygen in the air' the cause of the fire; it does not have the appropriate relation to the outcome. To amend this, Hart and Honoré propose a distinction within the set of necessary conditions, between "mere conditions" and "causes". (Hart & Honoré, p.32). The proper relation between a cause and an outcome is that a cause is "a *difference* to the normal course which accounts for the difference in outcome." (p.27, italics in the original) This seems sensible: "when things go wrong and we then ask for the cause, we ask this on the assumption that the environment persists unchanged, and something has 'made the difference'." (p.34) Indeed, it seems obvious that my match dropping is a difference to the normal course of events, while oxygen in the air is not.

But such a definition, one may note, simply transposes the theoretical burden to the term "normal course" - for what is deemed 'normal' then decides what counts as a cause. We need a further clarification for this definition to be useful. Hart and Honoré advocate to establish normality by common sense (such as in the case of the fire), but they highlight that the difficulty in doing so lies in the fact that the answer is *relative to the inquiry*, in two ways. Firstly, the same condition may be normal in one context and abnormal in the other. In a laboratory set-up where oxygen is deliberately supposed to be absent, the presence of it is precisely abnormal. Secondly, the same set of conditions may generate a different labelling of (ab)normality depending on the perspective taken:

So the wife of the man with the ulcerated stomach, who looks upon the parsnips as the cause of his indigestion, in asking what has given him indigestion is in fact asking: 'What has given this man in his condition indigestion when usually he gets by without it?' The doctor who gives the man's ulcerated condition as the cause [...] is concerned only with [...] deviations from the standard physical condition of

human beings. His question (in contrast with the wife's) is: 'What gave *this* man indigestion when other men do not get it?' (p.34, italics in the original)⁸

The wife and the doctor end up with a different cause because they both hold different things constant. The wife takes as constant the conditions of her husband's general life; she contrasts the current situation with other days, and then his diet is what stands out. The doctor, in contrast, takes as constant the physical conditions and behavior of the general man (his age); he contrasts the current situation with other people, and then his ulcers 'make the difference'. To be sure, the fact that causation has a relative element need not deter us. Relative is not the same as arbitrary: people with the same perspective, holding the same things constant, should still come to the same conclusion. The important insight is that statements on causation require us to specify the question we are answering.

As a final preliminary, we may wonder whether the difference between SI and interactional injustice is fundamental or merely a matter of degree. Conceptually I suspect they are separate, since we can find clear examples of their pure forms. An intentional murder by someone who just felt like it, acting regardless of any structural constraints, seems a paradigmatic case of an interactional injustice. Likewise, Sandy's story has been construed in such a way that no agent acts wrongly – the landlords act reasonably and according to the law, housing agents go out of their way to help, etc. I think such a scenario is plausible. However, it must be admitted that in practice the two injustices can complement and reinforce each other. People in a situation like Sandy's could *also* be harmed by an overly greedy or sexist landlord. In such a mixture, I consider the injustice to be (also) structural if the causal role of social structure is necessary and important for a full explanation (which I argue to be the case for Sandy). Due to space constraints I cannot elaborate further on the evaluative terms 'important' and 'full'; I assume that there are clear enough cases for these terms to be meaningful.

1.2.2 *The causal claim: structural factors as a causal explanation*

Let us now examine the role of structural factors in Sandy's story. What causal question about Sandy's situation are we trying to answer? As Haslanger remarks, "not all questions are equally good at probing the phenomenon in question to get at a good explanation." (2016, p.116) We need to ask the right one. I contend there are two options. The first is to ask why *Sandy* got in this situation, as opposed to other people; we hold constant the background conditions on which Sandy operates. The second option is to ask why Sandy got in this situation *together with many other people* - Young uses empirical facts to convincingly argue that it happens all over the US. (see e.g. p.44) The relation between the first and second question is that between *token* and *type*: Sandy's specific situation is a concrete instance of a pattern that applies on a much larger scale. (Sangiovanni, 2018, p.463/464; Hart & Honoré, 1959, p.31)⁹

What we do with a type-explanation, to start there, is wonder why so many people with different lives end up in the same situation. We hold constant individual choices, and investigate the *background conditions* of those choices as 'making the difference'. These background conditions, I argue, can only properly be understood as social structures. There are cultural expectations of mothers; the gender wage gap; inflation in the housing market because of buyers' high expectations; deficient public transport due to the cultural values attached to cars; segregation due to cultural preferences, creating 'nice' and 'less nice' neighborhoods; laissez-faire housing policy; and the list goes on. All these elements constrain agents, and are experienced as objective. Sandy, together with many others and regardless of her life choices, is *vulnerable* to homelessness. This vulnerability is an emergent property of large numbers of people acting a certain way, chasing their own goals. Fully understanding this

⁸ Though not specified in the original passage, let us assume for the sake of example that neither of them is in fact wrong. Suppose that both the eating of the parsnips and the stomach ulcer are relevant causes, and that both the wife and the doctor knew of both. Otherwise one could object that the wife would have given the same answer as the doctor, had she reasoned rationally with full information. And if that were the case, then the meaning of 'normal' would not have been relative to the perspective but simply a matter of rational deliberation.

⁹ In Hart and Honoré's mention of this distinction, they are quick to point out that the perspective of type-explanations mostly belongs to the sciences, while "ordinary life" - and presumably they have in mind law and philosophy - are more concerned with token-explanations. But this dichotomy is precisely what the turn to structures tries to break: it is also philosophy that can (and should) be concerned with type-explanations.

vulnerability means identifying as causes socio-structural effects, the interplay between schemas and resources.

A skeptical reader might point to a malicious and powerful actor as cause of these background conditions, rendering Sandy's situation an interactional injustice after all.¹⁰ For example, a government implementing bad policy for own gain could also explain why many different people suffer the same vulnerability. To this I reply that my structural analysis does not exclude malicious agents; I only claim that the explanation cannot be limited to them. I believe a malicious government is an insufficient explanation here, for two reasons. Firstly, notice how all the structural elements mentioned in the previous paragraph are not enough by themselves to create widespread vulnerability to homelessness: only combined do they generate this result. Bad policy only creates vulnerability in conjunction with certain cultural values and economic processes, which we cannot all ascribe to the government. Secondly, only focusing on a powerful agent means neglecting the structures in which *they* are embedded: governing bodies are also constrained by voters, lobbies and other governing bodies. Referring to malicious agents simply does not tell the whole story.

Let us then look at the second sort of explanation we can give for Sandy's situation, the token-explanation. This is where we hold steady the background conditions, and wonder why *Sandy specifically* ends up like this. A complete answer can consist of three elements. The first concerns her own choices in life: what is 'abnormal' and thus a cause is perhaps the fact that she did not go to college and left the father of her children. This is also important: regardless of structures, Sandy still has agency over her own situation. As Young observes, "part of what it means to respect people is to consider them agents of their own lives," and so we should Sandy. (p.27) The second option is to look for a malicious agent and interactional injustice, like a shady landlord. But again, the example is construed in such a way that there is none, which is a realistic scenario. Then the third option is to point to structural effects, also on the token-level. If social structure has influence on a large scale, causing many people to be in a certain situation, then this influence must also apply in most individual cases. The fact that a sensible life choice can lead to homelessness, such as leaving the father of your children (for a good reason we may assume), is ascribable to the society in which it takes place. Young notes by analogy that "a person's vision impairment impedes his ability to be a good civil engineer only in a society that fails to offer educational and employment opportunities that accommodate persons with differing physical abilities." (p.31) This underlines how personal circumstances are only positive or negative in relation to the environment in which they are located – and this environment is shaped by social structures. Thus also on the token-level, explaining Sandy's situation requires highlighting structural effects.

1.2.3 *The evaluative claim: wrongness*

Structural injustice so far can be phrased as "social processes put[ting] large groups of persons under systematic threat of domination or deprivation of the means to develop and exercise their capacities." (Young, p.52) But why is this an *injustice*, and not merely an instance of bad luck? One might argue that, precisely *because* a problem is caused by structures rather than by direct interactions between agents, it cannot be said to be wrong. It would rather be *unfortunate*, since wrongness can only be generated by human actions and not by something abstract as a structure. Such a position, however, ignores the fact that social structures exist by virtue of human actions. They *appear* as objective, but "a structure is an emergent property of social practices." (Aragon & Jaggard, 2018, p.442) Hence, they are human-made. This is what makes them different from (say) a meteor falling out of the sky, and what makes them interesting from the perspective of responsibility. I am not yet claiming anything here about the moral salience of the actions generating SI (that will happen in chapter three), I merely claim that structures do not *exclude* wrongness.

Rejecting the exclusion of wrongness is of course not the same as positively establishing wrongness. But the latter requires a substantive theory of justice, which I think is not necessary here: just as we can speak about responsibility for interactional injustices without defining what those are exactly, so too with responsibility for SI's. As Zheng notes, "a given theory of responsibility is in principle compatible with a range of different first-order ethical or political theories." (2019, p.113) Examples of SI's in later chapters include sweatshop labor and contract slavery in Thai brothels. I

¹⁰ I am grateful to colleagues Max Westhof and Ruben Tricoli for making this argument.

think it is uncontroversial to assume that most 'first-order theories' would consider those situations unjust.

1.3 Conclusion

In this chapter I argued that social structures can cause injustice, and therefore that the concept of structural injustice is a useful analytical tool. Social structures have been characterized as constraining in an option-shaping way; recursive and often self-reinforcing; experienced objectively; dynamic, in that they exist in the actions of masses of individuals; and mostly unintentional. Crossing over domains (economic, political, cultural), structures can create widespread positions of vulnerability to deprivation and domination. This causal role of structures became clear by conceptualizing causes as necessary conditions that 'make the difference' in the arising of such vulnerability. The role of structures becomes most apparent when we try to explain patterns of large-scale occurrences (types), but they also apply in particular cases (tokens). We contrasted such structural injustices with interactional injustices, in which there is a more direct relation between perpetrator and victim.

However, if we accept that social structure can cause injustice (instead of agents), we have to wonder what this means for responsibility. Who (if anyone) is responsible for SI, and on what grounds? These questions are the topic of the chapters to come.

Chapter Two

Conceptualizing responsibility for structural injustice

Now that we know how to think about structural injustices, we can start wondering what it means to be responsible for them. This is not an easy task. As Young notes, “the very judgment that there is injustice implies some kind of responsibility [but] when the injustice is structural, there is no clear culprit to blame...” (2011, p.95) A common strategy, then, is to examine the concept of responsibility itself: we can distinguish different meanings in the hope of finding one that applies in the structural context. Responsibility, however, can be rather elusive. The following passage illustrates the subtleties of its usage:

As captain of the ship, X was responsible for the safety of his passengers and crew. But on his last voyage he got drunk every night and was responsible for the loss of the ship with all aboard. [...] Throughout the voyage he behaved quite irresponsibly, and various incidents in his career showed that he was not a responsible person. [...] He is still alive and he is morally responsible for the deaths of many women and children. (Miller, 2007, p.82; Hart & Gardner, 2008, p.211)

This chapter is devoted to one goal: *clarifying what we mean with ‘responsible’* when we ask who is responsible for SI. To be sure, the aim is not to index *all* different possible meanings or types of responsibility; the focus lies specifically on being responsible for injustice.

In the first section I follow David Miller (2001, 2007) and distinguish two broad ideas of responsibility, under which a total of four specific ‘types’ of responsibility fall.¹¹ The goal of this section is to explain these four types and how they come apart. Idea 1 is concerned with the connection between agency and the arising of an outcome: it involves past-oriented concepts like blame and praise. The three responsibility types that fall under Idea 1 are causal, outcome and moral responsibility. Causal responsibility, however, is simply synonymous with causation: it does not have to involve agency, which makes it a bit out of place within Idea 1.¹² Because it functions as a necessary condition for the other two types, I still mention it under Idea 1, but I regard it as some distant conceptual cousin. Idea 2 revolves around remedial responsibility: it is concerned with ameliorating a bad state of affairs.

In the second section, making the connection with the previous chapter, I wonder which of these four types of responsibility apply in the structural context. I argue that some sort of causal responsibility can be established, but deny the same for outcome and moral responsibility, which are the core of Idea 1. Therefore, the most relevant and interesting conception of responsibility in the context of SI must be Idea 2: remedial responsibility. This is the main result of this chapter: *when we speak of responsibility for structural injustice, we intend remedial responsibility to combat the injustice*. The next chapter then deals with the question of what *grounds* such a responsibility. To foreshadow the discussion there: the crucial condition that generates remedial responsibility is causal responsibility. Even though it does not really belong to Idea 1, it carries some moral weight. We use the conceptual cousin from the first family as a criterion for allocating the second.

In claiming that we should not focus on blame but on tasks of redress, I echo Young. However, I deviate from her in a crucial respect: I reject her conceptual distinction between the Liability Model and the Social Connection Model, in favor of Miller’s. In the third section of the chapter I argue that Miller captures the essence of what Young tries to say, but without its problematic elements. In the fourth and fifth section, finally, I discuss two challenges to my position. The first consists of (two different) arguments claiming that Idea 1 *can* be extended to the context of SI, which I refute. The second challenge doubts whether Miller’s distinction is internally coherent in terms of temporal direction: is it at all possible to be forward-looking (remedially) responsible without any backward-looking (Idea 1) counterpart (such as guilt)? An objection of this kind has been prevalent in the

¹¹ To avoid confusion between the two general ideas of responsibility and the four types that fall under it, I sacrifice elegance and label the ideas ‘Idea 1’ and ‘Idea 2’.

¹² Miller does not address this issue.

literature (see Nussbaum, 2011; Barry, 2013, 2016, 2021) and deserves adequate attention. I believe this challenge rests on a misunderstanding. The sixth section concludes.

2.1 Two different ideas of responsibility, divided into four types

When we speak of responsibility for injustice simpliciter, we can have two different things in mind. The first is the broad idea of an agents' "involvement in the causal process that produces a bad state of affairs" (Abdel-Nour, 2018, p.16), Idea 1 for short. It is based on the belief that respecting someone as a person means acknowledging them as having some sort of agency over their own actions, and therefore treating them as the authors of their own life. The outcomes of their actions are 'theirs', and so the benefits and burdens of those outcomes should also be attributed to them. This is of course closely linked to ideas of desert, blame or praise, or punishment. By contrast, Idea 2 is that of remedial responsibility: "To be remedially responsible for a bad situation means to have a special obligation to put the bad situation right." (Miller, 2007, p.454) Such a responsibility rests on the acknowledgement of humans as finite beings, living in a world with injustice and disaster: the same agency that ties us to outcomes of our own doing might also be called upon to help others, just like we sometimes need to be helped ourselves. It belongs to the same category as "duty, obligation, and assigned or role-associated tasks." (Abdel-Nour, p.18)

To be clear, both ideas are often invoked together, like when a fraud is found guilty and has to compensate his victim along with an additional fine. But it is important to recognize that the two may come apart. An example of Idea 1 without Idea 2 is a child dropping a precious vase. We may attribute the shards on the floor to the child and even blame it for its clumsiness, but we often choose not to let it pay for the damages. An example of the second without the first is that of natural disaster or accident: in the child-in-the-pond scenario, surely the first passer-by acquires the responsibility to save it, even though the situation is in no way attributable to her. Besides them coming apart it is also important to recognize that, even if they are invoked together, one of the two is often given priority. Imprisoning a murderer may perhaps count as remedial responsibility according to some theories of punishment (on which I remain neutral), but the *main* endeavor revolves around blame and punishment. We first concern ourselves with Idea 1 and Idea 2 follows almost as a corollary. As Young observes, "practices of assigning responsibility in law and everyday moral life first try to locate 'who dunnit'" (p.95). By contrast, it is also possible to be primarily interested in remedial responsibility and merely use answers to 'who dunnit' as *criterion* for allocating it. This would be the case when I borrow my camera to a friend but it breaks in his possession. I may demand restitution of my broken equipment from him because he was responsible for it, but my primary concern is with the compensation. There need be no emphasis on fault or blame.

These general ideas already provide a useful framework, but Idea 1 is still an umbrella-term: Miller identifies three types of it (bringing the total of responsibility types to four, along with *remedial responsibility*). This deeper distinction is relevant because these types, as it will turn out, are not equally applicable in the context of SI. They are *causal*, *outcome* and *moral responsibility*, which I will discuss in turn.

"To say that an agent is causally responsibly for some state of affairs is to highlight the causal role played by the agent in the genesis of that state of affairs." (Miller, 2001, p.455) This is the type of responsibility we mean when we say 'the drought is responsible for the loss of my crops'; it is synonymous with causality. Thus any theory on causality can be used to identify causal responsibility – I will use Hart and Honoré's conception of causality as discussed in the previous chapter. To recap: something is causally responsible if it is a necessary condition for something happening, and if it deviates from normal, thus 'making the difference'. Furthermore, in principle the length and nature of the causal chain do not matter. If a causal analysis (à la Hart and Honoré) gives a certain result, this is where causal responsibility lies, regardless of how far away this cause is from the effect. "A's releasing a butterfly in China might in theory be identified as causally responsible for a hurricane in the Bahamas." (Miller, 2007, p.87)

Now Miller specifically remarks that agency is not necessary for causal responsibility (like in the example of a drought). But this seems odd, given that Idea 1, to which causal responsibility belongs, concerns connecting *agency* with an outcome. I think Miller still includes it in Idea 1 because causal

responsibility denotes a relation between an *event* and an outcome, which is related; and because it is a necessary condition for the other two types. This makes it relevantly related to them, but not strictly part of the same family – more like a conceptual cousin.

Continuing with the second type, outcome responsibility, “we want to know whether a particular agent can be credited or debited with a particular outcome.” (id.) We wonder whether the outcome can meaningfully be said to be ‘the agent’s’, as opposed to someone else’s or a mere happening. This is a relevant question when it comes to desert: we may wonder whether the merits and burdens of a certain outcome are justly distributed. Causal responsibility is a necessary condition, because one generally cannot be attributed the arising of an outcome without being causally involved in it.¹³ But additionally, the causal relation needs to be of the *right kind* to link an outcome to a person. This involves several conditions. Firstly, the act causing the outcome must be a genuine act of agency, which excludes bodily reflexes. Secondly, *here the length of the causal chain matters*. This makes sense: the longer it is the more shackles it has, each of which could change the direction of the chain (or break it altogether). The link between the agent and the outcome dissipates, so that the outcome cannot be said to be ‘the agent’s’ anymore. Thirdly, even if the agency is directly causal, it needs to reflect the skillset and possibilities of the particular agent. In the case of a positive outcome, this translates to the exclusion of sheer luck. Miller gives the example of an absolutely inexperienced golf player hitting a hole-in-one on an incredibly difficult course (2007, p.88). Even though the swing is an instance of agency, it would be odd to attribute the hole-in-one squarely to the player, since it is not in her set of capacities to do something like this. In the case of bad results we may demand that the outcome could reasonably have been avoided by the average person, again linking the outcome to the actual skillset of the agent. This is to account for consistently clumsy people: “even though, at the time of the events, he could not have averted [it...] we expect such a person to be aware of his shortcomings” and to take precautions to avoid the outcome. (id.)

It should be noted that some evaluative concepts have slipped into our conception of outcome responsibility. When is a causal chain too long, or a matter of sheer luck, and when is an outcome reasonably avoidable? I will not adjudicate boundaries on these scales: I assume that the extremes can be intuited clearly enough. All we need for our present purposes is the general conception and awareness of the existence of these variables.

The third type is moral responsibility: one must be outcome responsible, but additionally “the agent’s role in bringing about the outcome must be such that it leaves the agent liable to moral blame” or praise. (Miller, 2001, 456) What leaves an agent liable to such moral reactions is dependent on an external (moral) theory. As remarked by Zheng and quoted earlier, “a given theory of responsibility is in principle compatible with a range of different first-order ethical or political theories.” (2019, p.113) To illustrate how outcome and moral responsibility can come apart, suppose an ambulance crashes into another car while rushing to the hospital. The ambulance driver acted in full agency, had a direct causal link and the accident (let us suppose) could reasonably have been prevented, so she is outcome responsible and her insurance must pay for the damages. But given the job she was doing, I contend we would not blame her for the trashing of the car; she is not morally responsible for the damages.¹⁴

To sum, we have four different types of responsibility: causal, outcome, moral and remedial. The first three belong to a general idea of connecting a person to a specific outcome (of which causal responsibility is the odd-one-out). The fourth belongs to the general idea of assisting fellow humans in troublesome situations. In the next section I evaluate these responsibility types in the context of SI: which of these types is applicable? I will claim that causal responsibility can be somewhat established, but given that it is just causality, it is not really interesting qua responsibility type. Outcome and moral responsibility, on the other hand, are potentially salient but cannot be established for SI. We

¹³ As Young notes, an exception to this is complicity, in which an agent merely enables the action of another agent (p.96). See section 2.4.2 for a discussion on this.

¹⁴ Concerning the relation between moral and causal responsibility, Miller implies in an earlier paper (2001) that one can be morally responsible without being causally responsible. His example is a parent who pays no attention to his children playing, causing one of them to break the arm of the other. But using Hart and Honoré’s framework, I believe the parent’s negligence is precisely out of the ordinary and thus a cause: omissions can count as well. More generally, bracketing complicity, it seems problematic to be morally responsible for something without also featuring in the bringing about of that something. It appears that he dropped this stance in later work (2007).

should therefore aim our attention at remedial responsibility. In chapter three I argue that causal responsibility is the ground that generates it: even though it may not be that interesting as ‘responsibility’, it still carries some moral weight.

2.2 Evaluating the responsibility types in the context of structural injustice

Let us start with causal responsibility. In the previous chapter we established that social structure can be the cause of injustice, so social structures can have causal responsibility. But this alone is an unsatisfying result: we want to know what responsibility we as *agents* have. The strategy is this: if we can establish that agents are causally responsible for social structures, and social structures cause SI, then agents are causally responsible for SI. Recall that for causal responsibility the nature and length of the causal chain is irrelevant. From the perspective of causality it does not matter if there is a structure mediating between agent and injustice; causal responsibility is transitive.

In a way, we have established a causal relation between agents and structures already. Recall that social structures are emergent properties from the uncoordinated actions of large numbers of agents: they exist by virtue of the actions of people. This is precisely what distinguishes SI from problems like natural disasters: they are contingent and human-made. Resorting to Sandy’s situation again, we may ask ‘what has caused the existence of the economic structure of housing prices’, or ‘what has caused the existence of the cultural structure of demandingness on single moms’? The answer invariably is ‘the actions of millions of people’, thus establishing a structure.

But there are a few caveats. Firstly, under a basic *sine-qua-non*-test each *individual* contribution cannot be a cause of social structure, because it is not necessary for the outcome. We only know that *large groups of actions* cause social structure; it is not trivial how this translates to individuals. A solution to this might be to speak of *partial* causal responsibility, as Alvin Goldman (1999) does in the case of citizens voting. A single vote, he discusses, is never sufficient for an outcome, but also practically never necessary, as that would mean that subtracting one vote would change the outcome. But if we take this to mean that no single vote(r) has causal responsibility for the result, then there is no causal responsibility for the result at all, which is a paradoxical conclusion. Goldman proposes we see voting as a ‘vectorial causal system’, mimicking the interplay of mechanical forces in physics: “when [the outcome] in a vectorial causal system moves in a given direction, this is because the sum of the forces [...] are positive in that direction,” like in a tug-of-war. (p.210) Each vote is then a vector or force that adds to the total sum of the forces. If one votes for the winning party, one has contributed to the outcome and shares partial causal responsibility for it. What matters is thus not whether a contribution was necessary, but whether one contributed positively to the outcome that occurred. I argue that social structures can also be characterized as a vectorial causal system, and thus that it makes sense to also speak of partial causal responsibility for them.¹⁵

Secondly, we have seen that structures are recursive. They are contingent on the actions of people, but they constrain people’s action in turn. And the abstract cultural values of the structure influence the resource distribution, and the resource distribution shapes the cultural values in turn. This recursion breaks the dichotomy of simple (non-)existence: rather than the simple ‘being vs. not being’, existence is gradual, spreading out over time. Therefore causation in this context is different from the ‘normal’ causation in ‘my fist causes this dent’: it is rather that millions of agents cause the structures to *keep existing*. Accordingly, to speak with Young and others, it seems more appropriate to speak of agents reproducing structures than causing them. Combining this with the previous point, it can be said that agents *participate in the reproduction* of social structure. McKeown gives the example of English as a language: people speaking it are not causing English as a language, but with each usage legitimize and strengthen its existence. (2018, p.497) I still think this is a version of causal responsibility: we are still concerned with the relation between an event and an outcome, but over a longer period of time.

¹⁵ A different way to approach this first caveat is to say that the *group* of voters is causally responsible, and that responsibility is passed along to its members by virtue of its membership. But such an approach would require some social ontology on the concept of a group (membership), making Goldman’s approach more parsimonious.

It is worth noting that Young rejects the notion of causal responsibility for SI, because individual actions contribute “rather *indirectly*, collectively and cumulatively.” (p.96, emphasis mine) Simultaneously, however, she does acknowledge the salience of reproducing structures by participating in them – which seems paradoxical given that such reproduction is equally collective and cumulative. I suspect she worked with a different idea of causal responsibility, judging by her worry of indirectness: perhaps she had something like Miller’s *outcome* responsibility in mind, because for that the length and nature of the causal chain matter. In that case I would warmly agree with her. A single agent’s contribution is so far removed from the actual SI (causally seen) that the outcome cannot meaningfully be said to be the agent’s. The link between agent and SI is mediated by structures, which exist due to masses of people acting in a certain way across different domains (economic, cultural, social) – instead of a causal chain we have a massive causal web. Outcome responsibility for SI seems misplaced.

But if we cannot establish outcome responsibility, then the same goes for moral responsibility, as the former is a necessary condition for the latter. This is what I take Young to mean when she argues that no fault or blame is appropriate for SI. To be sure, blame or guilt can still be appropriate for an *interactional injustice* that accompanies the SI. A malicious landlord unlawfully expelling Sandy could still be blamed for this. I argue that blame is unfair to those regular agents that mind their own business but very indirectly contribute to the structures with their everyday actions.

So from Idea 1, we have established that it is appropriate to speak of partial causal responsibility in the context of SI, but of outcome or moral responsibility. This is not much, especially considering that causal responsibility is equivalent to dry causation, bearing no relation to the agency of the people involved (as with outcome responsibility). I therefore propose we turn to remedial responsibility as our prime conception of responsibility – we put the emphasis on Idea 2 instead of Idea 1. We are first and foremost concerned with combatting the SI, and start thinking of Idea 1 only in second place. I believe this makes sense given the ruminations in this section, and the massive harm that SI’s can generate: if all alleged instances of SI are indeed that, it is widespread and pervasive, affecting large masses of people who can suffer severely from it. This makes it legitimate, if not imperative, to question who, if anybody, is tasked with ameliorating this. Answering this question means formulating a criterion or *ground*; this will be the topic of the next chapter. The current chapter has, so far, only argued that it is *this* conception of responsibility that we are after.

2.3 Young’s distinction: The Liability Model and the Social Connection Model

So far my analysis has differed from Young in a crucial respect: instead of Miller’s terminology she developed her own responsibility distinction. Her Liability Model (LM) and Social Connection Model (SCM) somewhat resembles Miller’s Idea 1 and Idea 2 respectively. In this section I want to briefly argue why I favor Miller’s: it captures Young’s intention more accurately.¹⁶

Young describes the LM as the “most common model of assigning responsibility [deriving] from legal reasoning to find guilt or fault or harm.” (p.97) Causal responsibility is a necessary condition, as well as some variables of agency such as voluntariness and intention. The LM is backward-looking in a temporal sense, as it seeks to attribute blame or guilt for something that has already occurred. She includes “all such practices of assigning responsibility [...] that seek to identify liable parties for the purpose of sanctioning, punishing, or exacting compensation or redress.” (p.98)

The SCM contrasts sharply with this. She describes it in five features (p.105-113). Firstly, it is ‘not isolating.’ Instead of aiming to identify a few perpetrators, it seeks to identify “all who dwell within the structures [who have to] take responsibility for remedying injustice they cause”. (p.105) As a corollary of this, secondly, responsibility under the SCM is essentially shared with others. Thirdly, this makes the actions it prescribes *collective* in nature. Fourthly, the SCM focuses on *background conditions*, on ‘the normal’, while the LM focusses on ‘abnormal’ deviations from the background conditions.

¹⁶ There are other, comparable distinctions available in the literature. See Zheng (2019) for ‘attributability’ vs. ‘accountability’, and Beck (2020) and Lu (2017) for ‘interactional responsibility’ vs. ‘structural responsibility’. I lack space to describe in detail my arguments against those distinctions, but they are similar to those against Young: the others commit themselves to burdensome conceptual redundancies, or emphasize irrelevant distinctions. Miller’s account captures the relevant distinction *in nuce*.

Fifthly, it is forward-looking instead of backward-looking. Next to these characteristics, Young supplies four “parameters of reasoning” that help establish how much one is responsible for under the SCM, depending on the social position one occupies. (p.142-151) They are: power (being in a position to change it), privilege (benefitting from unjust structures), interest (being benefitted most by changes in the structure) and collective ability (having access to existing networks or institutions that can make a difference).

Why should we favor Miller’s distinction? Because I think it is clearer and more concise: it points in the same direction but without the extra theoretical burdens that Young imports. Notice that the SCM’s five features are not only about what responsibility means (about types), but also about how to discharge it (collectively) and the type of injustice that it concerns (background conditions). These matters are logically independent from responsibility types: a certain type can in theory have multiple ways of discharging it and can be generated by different injustices. Miller’s terminology allows this: remedial responsibility can for example be generated both by SI and interactional injustice. Young takes multiple distinctions (collective/individual, background/deviations from background) and treats them as if they are connected, bundling the extremities into two models – but the distinctions are in fact orthogonal.

Moreover, the features that *do* concern responsibility type actually blur the distinction rather than clarify it, as they can apply to both the LM and SCM (realistically conceived). (Beck, 2020, p.10-12) Firstly, the LM and SCM can be equally isolating. Young argues that the LM has “the aim [...] to assign liability to particular agents as isolatable and identifiable”. (p.105) But the SCM does the same: it formulates a criterion for responsibility allocation (“dwelling in the structure”, or ‘being connected to it’) and identifies agents who match these criteria. I fail to see how this is different. It cannot simply be that *more* people fit the criterion under SCM than under the LM, because that is a contingent fact, independent of the conceptual distinction. Secondly, responsibility under both LM and SCM can be shared. The LM also applies to collective agents and it can identify multiple agents at the same time. Young may have believed otherwise because, as it has been pointed out, she applied the basic *sine-qua-non*-test for establishing causal responsibility: “it should be the case that if I had not done what I did, then the harm would not have occurred, or would not have occurred to the same extent.” (p.101) Under such a conception it may indeed be difficult to find shared fault, because only the decisive contributions matter (which is a problem of overdetermination). But the *sine-qua-non*-test is rejected or amended by most scholars because it generates counterintuitive results (see section 1.2.1). Under an improved conception of causality all those who contribute to the injustice can be liable, making the LM shared as well. Thirdly, responsibility under both the LM and the SCM can prescribe collective action. Under the LM a group of convicts might be tasked with collectively redressing their crime, such as doing community service together.

I believe Young had a very similar distinction in mind as Miller: what she saw as the LM primarily concerns the relation between one’s agency and a certain outcome (involving elements of guilt and blame), and her SCM primarily concerns remedying the injustice that is occurring. But Miller’s version takes *only* this, leaving theoretically independent elements aside, and preserving the clarity of the distinction.

2.4 Counterarguments extending moral responsibility to structural injustice

This section considers two arguments claiming that we *can* extend Idea 1 (outcome and moral responsibility) to the context of SI, which I both refute. The first, Parfit’s Harmless Torturers example, fails mainly because it misconceives the causal nature of SI (exemplifying the relevance of the previous chapter). The second is an argument of complicity.

2.4.1 The Harmless Torturers

Derek Parfit (1986) has offered the thought experiment of the Harmless Torturers (HM), with which he tries to persuade us that “It is not enough [for moral responsibility] to ask, ‘Will my act harm other people?’ [instead] I should ask, ‘Will my act be one of a set of acts that will together harm other people?’” (p.86) Here I present Sangiovanni’s version of the experiment, who slightly modified the

original to better resemble SI.¹⁷ (2018, p.469) Suppose a thousand people have in their home a light switch that is connected to one torture machine that operates on another person. If all 1000 are switched on, a current runs through the machine which inflicts unbearable pain to the person on the other end – if all are switched off, the machine is off as well. Any single light switch at most inflicts a barely perceptible change, and most of them even none. Each person is aware of his involvement in the total current but switches the light on only for the sake of the light. Each household, furthermore, would have to pay a significant amount of money to rewire their electricity, and any specific quantum of current running through the machine is epistemically not traceable to any specific household (due to some form of anonymization).

Now we have a situation in which many people together create severe harm but do not do so individually; separate contributions cannot be traced to specific individuals; alternatives are available but costly (resembling the constraining effect of structures); and no individual intends to harm the victim, as they all enact their own plans in life. The argument then goes like this: i) the HM-case is analogous to SI; ii) with HM we think moral responsibility applies; iii) therefore this also applies to SI.

Sangiovanni argues for ii) in an elegant, almost mathematical way, which is worth the read but which I will not recite here, because my focus is on i). I submit that HM is not really analogous to SI, echoing two reasons given by McKeown. (2018, p.498-499) Firstly, in the HM case the harm is cumulative but *direct*. A flick of a light switch directly influences the outcome of torture. With SI we instead have a causal chain that is not direct and often not even linear. Both Young and McKeown illustrate this using the example of sweatshop labor in low-wage countries, where economically vulnerable workers are being exploited under bad working conditions. The existence of sweatshops is causally possible thanks to all parties in the production chain, as each demands the best results for the lowest prices from the party he has to deal with. This includes the consumer and the sweatshop-boss, but also a large number of subcontractors in between. Moreover, each link in the chain has its own normalized cultural values (consumers buying fashion, subcontractors maximizing profit); the workers are already vulnerable to such exploitation because of the bad social position in their country; governments generally do not have adequate legislation in place because they want to attract foreign investment; and the list goes on. In short, SI is causally multi-faceted, whereas the HM-case is not. This difference between a direct chain and a complicated causal web is relevant because of what we remarked earlier on outcome responsibility: the longer and more complicated the causal chain, the weaker the link between an agent's action to a certain outcome. It makes less sense to attribute an outcome to the agent.¹⁸

The second reason for rejecting the analogy is that it is not at all clear what the alternatives are in the case of SI. With the HM this is clear: either live in the dark or rewire the electricity, both of which would clearly solve the problem. With SI, however, solving it is far less clear-cut. Considering sweatshops again, simply abstaining from buying clothing produced in sweatshops may not help, because the stores might then simply move to a region (with perhaps even more industrious consumers). And if every consumer everywhere abstained then the vulnerable sweatshop workers would simply be jobless, possibly increasing their malaise. Instead, "this is less of a 'do or don't situation' and more of an 'act in a certain way' situation." (id.) Perhaps one needs to protest, or vote for certain policy makers, or spread awareness of the problem, or donate to an NGO, etc. How to change the structures one lives in is a question shrouded in empirical uncertainty. Again, this difference is morally relevant, because this uncertainty obfuscates the boundary between a 'right' and 'wrong' action. Is buying sweatshop-produced clothing right or wrong? Without a clear answer moral responsibility seems misplaced.

2.4.2 Complicity

The second type of extension of moral responsibility to SI is based on complicity. Although complicity is a difficult term to define and is also often confused with an entire family of cousin-like terms, it is widely held that complicity can generate moral responsibility (blame, guilt) without causal or outcome responsibility as a necessary condition. (Lepora & Goodin, 2013). This is promising, because

¹⁷ See Sangiovanni, 2018, footnote 18 for a full expression of how Parfit's example has been modified.

¹⁸ McKeown presents a third reason for rejecting the analogy: with SI there exist intermediary agents that can "change the nature of the harm and could alleviate it." (2018, p.499) This is true, but I fail to see how this is distinct from the first reason, as intermediary agents simply increase the length of the causal chain.

the lack of outcome responsibility was the main reason there was no moral responsibility in the case of SI. Thus, if we could find a suitable account of complicity, we could avoid all the conceptual gymnastics on responsibility in this chapter and simply apply that. I follow Young in regarding Christopher Kutz' treatment of complicity as the most promising out there. (Young, p.101-104) I will argue it is unconvincing in the structural context, echoing an argument by Sangiovanni. (2018, p.470-471)

Kutz argues in two steps: the first is that complicity, which generates moral responsibility, consists of participation in a collective project. His main example is that of the bombing of Dresden by the Allies in WWII: he believes (rightfully, I think) that even the bombers that intended to drop their load but failed to do so (say, because of a technical error), thus lacking causal responsibility, were morally responsible for the bombing. He explains this by them sharing 'participatory intent': they considered themselves as contributing to a collective goal (the bombing). The second step is to extend this to SI. Admittedly, people participating in SI generally lack such a collective project. People buying sweatshop-produced clothing do not share a common goal of producing sweatshop labor conditions. But, he claims, they can still be seen as 'sharing a way of life': they have similar beliefs, desires, values, expectations. If those people in those same structures then start *identifying* with this way of life, this creates the character of a collective project:

Individual agents are, broadly speaking, participants in a shared venture that does harm, and so are inclusively accountable for the unintended consequences of what they do together. Call this a *quasi-participatory* basis of accountability, 'quasi' because there is no specific project to which individuals contribute. (Kutz, 2000, p.186)

So even though there is no specific project, there still is a 'shared venture' by virtue of identifying with a shared way of life.

I think we have two reasons to reject this line of argument. To start, we may wonder whether people who participate in SI really share a way of life. Both a hippie and a corporate lawyer can buy sweatshop-produced clothing. But more importantly, as Sangiovanni rightfully remarks: why does *identification* with a shared way of life suddenly make it akin to a collective project? A collective project is, by Kutz' own lights, characterized by an intention to participate in achieving a certain goal. It is unclear how identification with a certain way of life should create such participatory intent, given that there still is no project or goal to obtain. If I identified myself as a sweatshop-produced clothing buyer, this does not make me have an intention of achieving sweatshop labor together with other sweatshop-produced clothing buyers - and even if I did form such an intention, such a project still would not exist. There is no coordination or goal that has been decided on by the collective. The bridge between a shared way of life and a collective project seems absent.

2.5 Incoherence in time-dimension

There is a final challenge to be discussed, that has been voiced frequently. (To name a few: Nussbaum, 2011; Barry, 2013, 2016, 2021; Digeser, 2018) It has consistently been directed at Young's distinction between the Liability Model and Social Connection Model, but it can be charged against Miller's as well, as both distinguish between temporal directions. I first present the objection along with my response to it, which differs from Young's view. I then discuss two new questions that my answer raises, and in doing so I also address why I think my solution is preferable to Young's.

2.5.1 The challenge

Miller's Idea 1 is primarily backward-looking: it is used in identifying responsibility for outcomes that have already arisen. Idea 2, by contrast, is mostly forward-looking: it focuses on what must be done in the future.¹⁹ But if we say that only the latter matters in the context of SI, then we have to wonder

¹⁹ For both ideas I say 'primarily' because this distinction is not strictly so. It is possible to use Idea 1 concerning the future: 'if you do that, you will be guilty of it'. Likewise, Idea 2 can use as allocation criterion a responsibility type from Idea 1 (as will be established in chapter three: 'you must remedy this because you have been causally

what happens if someone *fails* to shoulder that responsibility. We are faced with a dilemma. On the first horn, if we say one becomes blameworthy for such failure, a sense of backward-looking-ness intrudes on the forward-looking-ness. Say Alex is not blameworthy for Sandy's SI to begin with, but has remedial responsibility for it, as I have claimed so far. But say furthermore that he has not acted on this responsibility at all. If this makes him blameworthy, then from the moment SI started existing he has been blameworthy and non-blameworthy at the same time, which seems paradoxical (shortly I will argue that it is not).²⁰ According to the second horn of the dilemma, if we say that failure does not make him blameworthy, we preserve the distinction but we give anyone "a free pass indefinitely, since no task they have failed to shoulder ever goes into the debit or guilt side of their ledger." (Nussbaum, 2011, xxi) We have then ended up, so the argument goes, with something that is not 'responsibility' anymore, because the bearer does not have to answer to anyone (as the literal, etymological meaning of the word 'responsibility' prescribes). There is no mechanism of *holding* one responsible anymore, which seems crucial. Since both horns of the dilemma create significant problems, the retrospective/prospective distinction should be abandoned.²¹

From this point I depart from Young. Where she bites the bullet on the second horn, I instead focus on the first horn of the dilemma and argue this argument is based on a misunderstanding. (Young, p.113-122) It must be recognized what one is actually blameworthy *for*. Consider the famous case of the child in the pond. When Alex first sees the child ($t = 0$) he is not blameworthy for it drowning, as it simply fell in by accident. Alex' agency is not connected to that outcome. Imagine now that Alex, in full agency, decides to save his expensive clothes and not the child. Then from that point on (a few moments later, say $t = t'$) he *is* blameworthy for it drowning: his agency, which now includes an omission, has become connected in a relevant way to the outcome. He became outcome responsible for something he did not cause but let happen. But now suppose, additionally, that someone else arrives at the scene (right after Alex' decision not to help) who *does* rescue the child. This way the bad outcome is avoided and Alex cannot be blamed for the child drowning (as it did not). But I think we say he is blameworthy nonetheless, for acting immorally and *not shouldering his responsibility*, regardless of the outcome. What this thought experiment shows, is that blame for an outcome can be separated from blame for the failure to act responsibly.

The first horn of the dilemma claimed that, if we say Alex is blameworthy for not shouldering his responsibility for SI, he has been blameworthy and non-blameworthy at the same time. But this is not a problem, as these are two different instances of blame: he has always been *non-blameworthy for the outcome*, but becomes *blameworthy for failing his responsibility*, and these two can co-exist. Notice how SI is dissimilar to the child in the pond: refraining from action does not make one outcome responsible. The reason that Alex is not blameworthy for SI in the first place is that his agential contribution is causally so far away from the outcome that its influence is almost negligible. But if he then fails this remedial responsibility, this lack of influence does not change: he is still not outcome responsible, regardless of his omission. Thus we have two types of blame, and the first horn is no horn at all.

2.5.2 Two new questions

So far we have avoided the objection against a responsibility distinction in temporal direction. But the solution comes at a cost, raising two new questions that need addressing. The first is this: I have claimed that acting on one's remedial responsibility for SI has an imperceptible influence on the outcome. But why should I care about it then? Why should one even have remedial responsibility if the actions it prescribes have negligible impact anyway? I answer this along the lines of Goodin and Barry (2021), who propose to shift away from a consequentialist framework: acting on one's remedial

involved'). But these nuances do not matter for the objection raised; in general the backward/forward distinction holds.

²⁰ It has also been suggested that one can become blameworthy for failing one's remedial responsibility only after realizing one has it. This epistemic condition is an interesting topic that deserves more attention than I can give it here. See e.g. Lu (2018) and Digereser (2018) engaging each other on this topic.

²¹ So far I have spoken about blameworthiness, but the argument has been given in a variety of terms (among others, Nussbaum speaks of guilt and Barry of blame). In their response to Young this makes sense, as both terms are run together in her Liability Model. But in general, there is a family of terms like these (guilt, shame, fault, liability, culpability, moral responsibility, etc.) and it is an open question what their relations are. I have chosen blame here because it is the paradigmatic reactive attitude.

responsibility does not aim at obtaining a result (though it could do so if done in large numbers), but it rather means “taking a stand” against injustice, which is more a “virtuous or deontologically worthy gesture,”: it is “the right thing to do.” (p.346) Good consequences *may* ensue if many people do the same, but this is not the reason one should do it. They, too, mention the paradigm example of voting: “one can [...] vote for ‘more justice’ out of a desire to ‘take a moral stand’, even knowing that one’s own vote is extremely unlikely to be causally decisive.” (p.347) I think this is a compelling argument for why remedial responsibility for SI matters.

The second question that my solution to the initial objection raises (as voiced by Young herself) is the following. If one can be blameworthy for failing their remedial responsibility, is one then not blameworthy *always*? As Zheng points out, the world is filled with a “sheer enormity” of SI. (2019, p.120) It influences many spheres of human life (housing, work, climate, gender, etc.). As “we are finite beings subject to limitations of time, space, and energy”, we simply cannot take a stand in *all* of them. (Zheng, 2021, p.503) Moreover, as argued against the Harmless Torturers, it is not always clear which actions are the right ones in taking a stand. Simply not buying clothes produced in sweatshops anymore might not be the way to go; political engagement might also be necessary, which makes taking a stand even more demanding. We are destined to come short in our efforts at taking a stand. A theory of responsibility that makes us blameworthy almost by simply existing in the modern world seems unfair.

I submit that, even though blame is possible, it will not be applicable in many situations. To start, it seems to me that a person who does *everything* in their power cannot be blamed – failure is always relative to the agent’s possibilities. But what about people who do *something*? There is a grave evaluative uncertainty here: it is very difficult to ‘measure’ when one has done enough to have ‘taken a stand’. Without such a threshold for success and failure, establishing appropriate blame becomes problematic. Perhaps the only clear-cut case of appropriate blame is when someone does not want to take a stand at all. If one, in full agency, decides not to act on one’s remedial responsibility, there is no evaluative uncertainty and one is clearly blameworthy. This is the reason I prefer my version over Young’s: I believe this is the bare minimum of blame we can and should allow, which Young cannot. Moreover, considering blame in general, it is widely acknowledged that one of the variables determining whether blame is appropriate is the standing of the blamer, in particular whether the blamer is hypocritical. (Tognazzini & Coates, 2021, section 2.3) Almost all theories on blame consider hypocritical blaming to be inappropriate. But given the pervasiveness of everyone’s involvement in SI, any instance of blaming in this context is almost immediately hypocritical to some degree, and therefore inappropriate. Taking these two remarks together, allowing blameworthiness into the picture of remedial responsibility for SI does not result in unreasonably excessive blaming.²²

To summarize this section: the individual agent is not and cannot become blameworthy for the existence of SI, but can only become blameworthy for failing to act on one’s remedial responsibility. One should act on it because it is ‘the right thing to do’. Blame is only appropriate, however, in the case of not doing anything at all; as long as one at least tries, evaluative uncertainty and hypocrisy make blame inappropriate.

2.6 Conclusion

In this chapter I elaborated two general ideas of responsibility: one concerned with the connection between the arising of some outcome and agency (Idea 1), and one concerned with the task of redressing a bad state of affairs (Idea 2). Of Idea 1, I have argued that only partial causal responsibility makes sense in the context of SI (which is synonymous with causation and not that interesting as ‘responsibility’), not outcome and moral responsibility. Therefore, if we want to establish responsibility for SI, we should look at Idea 2: remedial responsibility. While this result is fruitful, it is

²² There remains the interesting case of an agent acting on her responsibility, but very minimally, so that it is obvious and reasonable to suggest there is a lot of room for improvement in that agent’s behavior. I have only space for a brief outlook on this. It has been suggested by Young (who did not elaborate on it further) and by Zheng (who did) that in such case *criticism* is warranted as opposed to blame. Zheng (2021) proposes a distinction between summative criticism which ‘keeps score’ (like blame), and formative criticism which aims at improvement (like feedback).

still incomplete: we still do not know *who* has such a responsibility. The next chapter establishes criteria that answer that question.

I have also argued that Young's distinction between the Liability Model and the Social Connection Model fails because it blurs the distinction it tries to make – a distinction that is captured *in nuce* by Miller. Furthermore, I have rejected two possible extensions of moral responsibility to the structural context: that of the Harmless Torturers, and that of complicity. Lastly, I have rejected the notion that Miller's distinction is incoherent: one can only be blamed for refusing to act on one's remedial responsibility, not for a lack of outcome. Such blame is only appropriate if one does not even try.

Chapter Three

Allocating remedial responsibility

In the previous chapter we established that we are interested in remedial responsibility, i.e. the task of redressing structural injustice. But any responsibility type is not yet useful as a conception alone: it needs an additional theory to establish *who*, in fact, has that responsibility. This is what I call, after Zheng, a *ground* of responsibility (or allocation key): “what are the conditions in virtue of which someone bears or does not bear responsibility?” (2019, p.114) Causal responsibility needs a theory of causality; outcome responsibility, as mentioned in section 2.1, needs among others an account of agency; moral responsibility a substantive moral theory about blame or guilt. So too with remedial responsibility: it needs an account of what normatively salient aspect about a situation generates remedial responsibility. This heavily depends on the context: in the case of the child in the pond, for example, the relevant criterion is the capacity to help the child (being in the right place at the right time, being able to swim) - but in other cases other criteria might be relevant (like being the family of the victim, or having agreed to it beforehand by means of a contract, etc.). So in this chapter, we wonder what such relevant criteria are *specifically for the context of structural injustice*; we want criteria that honor the unique nature origins of SI.

In the first section I answer this in line with how Young is commonly interpreted: the relevant normative factor in the structural context is *participation* in the harmful structures - in other words, as established in section 2.2, partial causal responsibility. The relevant criterion for Idea 2 happens to be the distant cousin of Idea 1. Although we already discussed causal responsibility qua responsibility type, we have not yet connected it to remedial responsibility - it is not obvious why participation would justify claims of redress. In fact, most theorists do not elaborate extensively on this, as they seem more pre-occupied with excluding blame (like Young). I want to establish this link explicitly by elaborating on the intuition that any involvement in a harm, regardless of agency, justifies some normative response.

In the second section I consider Mattias Gunnemyr’s (2020) criticism on this position. His arguments stress the relevance of an alternative ground for remedial responsibility: the capacity to solve the problem. I respond to this in two ways. On the one hand, I accept some valuable suggestions: participation in SI generates remedial responsibility only *pro tanto* (and so it must be weighed against other responsibilities, as opposed to being *all things considered*), and it can even be *excused* by the circumstances. On the other hand I argue that capacity cannot replace participation as the relevant ground. Due to the nature of SI, agents with the capacity to solve it are highly unlikely to exist. And even if they do, capacity generates responsibility by virtue of some general humanitarian principle, which is unrelated to the structural context and therefore less interesting here.

It should be noted that, besides causal responsibility and capacity, other possible grounds for remedial responsibility can be found in the literature (see e.g. Miller, 2007). However, these have been offered as possible grounds in general, without consideration of context. Trying to apply them to the structural context specifically, it almost immediately becomes clear that they would not suffice. Outcome and moral responsibility, for example, could in general also be relevant grounds but we have already established in the previous chapter that they cannot be established for SI. I therefore choose to omit a discussion on these alternatives, for the sake of brevity.

In the third section, finally, I address two worries of overdemandingness that the choice of partial causal responsibility (reproduction) as the relevant ground might evoke. I argue that victims, in some circumstances, indeed also bear some remedial responsibility for the injustice they are the victim of. But I also claim that responsibility can come in degrees, and that responsibility leaves more to the discretion of the agent than a duty. The fourth section concludes.

3.1 Reproducing structural injustice as relevant ground for remedial responsibility

In section 2.2 we established that all agents who reproduce SI have partial causal responsibility. Here I want to argue that, even though this is simply a form of causality, it has normative significance and as

such generates remedial responsibility. Such a conclusion is in line with Young and some of her commentators:²³

Our responsibility derives from belonging together with others in a system of interdependent processes of cooperation and competition through which we seek benefits and aim to realize projects. Within these processes, each of us expects justice toward ourselves, and others can legitimately make claims of justice on us. All who dwell within the structures must take responsibility for remedying injustices they cause, though none is specifically liable for the harm in a legal sense. (p.105)

However, not many have made explicit *why* this is the salient moral factor we are looking for. We need a rationale for why participation matters for remedial responsibility. I want to offer a few considerations.

Firstly, we are looking for normatively salient factors in the specific context of SI. Therefore it makes sense to ground remedial responsibility in what characterizes SI: the fact that it is caused by structures. But structures, as we have seen, are an emergent property of the actions of masses of individuals. On the one hand they differ from natural disaster in that they exist in virtue of actions: they are contingent and man-made. On the other hand, they differ from interactional injustice in that the link from action to injustice is mediated by the structure (so that outcome and moral responsibility are out of the question). Even though agency is too far away from the SI (in a causal sense), there still is *some* agency involved – buying sweatshop-produced clothing is no accidental happening. Given that this is the unique nature of SI, it makes sense to see this involvement as the relevant criterion.

Secondly, as mentioned by Young in the quote above, we exercise this agency in everyday decisions in pursuit of our own projects and conceptions of the good life. In these endeavors we expect justice be handed to us. But if we consider that those same agentic actions, though not wrong per se, result in injustice elsewhere, causing other people to not be able to pursue *their* goals and dreams, that should count for something from the perspective of fairness. Participants in SI can and should choose to take a stand against the limitations that structures, through their actions, impose on others.

Thirdly, I build on the intuition that any involvement in the arising of a bad state of affairs, regardless of the amount of agency involved, bears some normative salience. I believe this intuition is elegantly captured by an example of Bernard Williams (1981, p.28) that is also used by Abdel-Nour (2018, p.17) in the context of SI. Williams speaks of a lorry driver who, through no fault of his own but because of some freaky accident, runs over a child. We cannot say the lorry driver is outcome responsible, because his agency is not connected to the outcome (let us assume that the average person could not reasonably have avoided the accident). Accordingly, he is also not morally responsible. We may even doubt whether he is causally responsible: that depends on the causal analysis of the accident that happened – but what is clear is that he was involved in its arising. As Williams notes, this is what makes him *different from a spectator* happening to be around. The driver will most likely feel some sort of regret for what happened, along with other mental states or actions that go with this – reaching out to the parents, feeling the need to support them, etc. People will likely try to comfort the driver, “to move [him] from this state of feeling [regret] to something more like the place of a spectator.” (Williams, p.28) But as he remarks, “it is important that this is seen as something that *should need to be done*, and indeed some doubt would be felt about a driver who too blandly or readily moved to that position.” (id., emphasis mine) The example shows that the lorry driver’s involvement, although not his fault, warrants and even prescribes certain moral responses. In the case of SI, in which one can be involved while only partially causally responsible, I believe the right moral response is the bearing of remedial responsibility: trying to take a stand against the injustice that is occurring.

²³ For example, Aragon and Jaggar (2018) also select reproduction of SI as the salient moral feature, but call it ‘structural complicity’: “people are structurally complicit when they exercise their agency in ways that reinforce the unjust social structures in which they participate, regardless of their conscious intentions.” (p.449) I choose not to adopt their terminology because the term ‘complicity’ evokes associations with moral responsibility (as discussed in section 2.4.2).

3.2 Arguments from capacity to reject participation

Mattias Gunnemyr (2020) has contradicted my analysis so far, arguing that participation is neither necessary nor sufficient for remedial responsibility. In this section I consider his arguments, with the aim to do multiple things at once: clarify my own account in some respects, strengthen it with some suggestions of his, and consider the most promising alternative to participation as the relevant ground: *the capacity to solve the injustice*.

To start with why participation is not necessary, Gunnemyr argues that agents might exist who do not reproduce structures but of whom we would still say they bear remedial responsibility, because they have the capacity to solve them. He uses the stark example given by McKeown of contract slavery of underage girls in Thai brothels.²⁴ Such injustice is possible, among other things, because “local economic conditions following rapid industrialization have given working-class Thai male laborers greater disposable income, social attitudes permit prostitution, and there is an assumption that children owe a debt to their parents.” (p.573) He wants us to suppose, by way of a thought experiment, that there exist agents who are not involved in these processes, but who “without significant cost [...] might, for instance, be able to implement a combination of measures reducing the demand for prostitution, changing some crucial social norms and improving the economic situation for the relevant groups.” (id.) Surely, he argues, these agents would have the responsibility to implement such measures.

My response to this is twofold. Firstly, given a genuinely structural injustice, it seems unlikely that there exist agents, even in principle, who have the capacity to solve it ‘without significant cost’. Assuming the opposite seems to me to misconceive the causal nature of SI. As discussed, structures are self-reinforcing, consist of millions of agents acting in a certain manner, and often feature complex interactions, resulting in substantial epistemic uncertainty on the best way to change a structure. Surely some agents are more powerful than others (like governments perhaps, or multinational corporations) and they have a much larger capacity to contribute to change, but even they have no simple solution ready. In the brothel example, simply banning brothels might make them go underground, ‘crucial social norms’ are definitely not easily changed, and any improvement of the economic condition of the relevant groups will probably not be ‘without significant cost’. A child-in-the-pond-like scenario just seems highly unrealistic with respect to SI.

But let us for the sake of argument suppose that Gunnemyr’s thought experiment could be realistic. Then secondly, I accept the conclusion that substantial capacity generates remedial responsibility, but point out that this is due to a general humanitarian principle of aid towards fellow humans (perhaps even towards fellow beings). Stephanie Collins (2018) discusses a version of such a principle, which she calls the Assistance Principle: “if you can fulfil important interests, at not too high a cost, then you have a moral duty to do so.” (p.354) I do not deny that such a principle exists, but my inquiry here focusses on normatively significant aspects in the structural context specifically. Therefore such a humanitarian principle is not very interesting. It certainly does not retract any significance from participation as a ground – the arguments made in the previous section still stand.

Gunnemyr also argues that reproduction of SI is not even sufficient, because there might be agents who indeed participate but of whom we would say they do not bear remedial responsibility. His main example is that of the victims of SI, who also participate (sweatshops exist thanks to sweatshop workers) but sometimes lack reasonable capacity to act on remedial responsibility. On the teenagers in Thai brothels he remarks euphemistically: “the opportunities [they] have to broadcast their situation (or unionize) are clearly very limited.” (p.578) Allocating them remedial responsibility seems overly demanding, and any instance of blame or criticism at their address seems misplaced. Therefore reproduction of SI seems not to be a sufficient condition.

I agree with the argument that in practice not *all* participants in SI bear remedial responsibility, depending on the circumstances. But I do not think that this means we have to abandon participation as the normatively salient point. We can simply apply the two modifications that Gunnemyr himself suggests to solve the issue. The first is that the remedial responsibility generated by one’s participation in SI is *pro tanto* instead of *all things considered*: it is a responsibility that has to be weighed against

²⁴ McKeown (2015) indeed uses that example to illustrate her case, but I have my doubts about how structural that injustice is (as opposed to more interactional). Due to space constraints and for the sake of argument I give it the benefit of the doubt.

other responsibilities and may be overridden if others are stronger. This modification seems appropriate for cases like a big family in a western country barely scraping by: by participating in several SI's (like buying sweatshop-produced clothing) they make themselves remedially responsible, but this responsibility has to be weighed against responsibilities like clothing one's children, putting food on the table, etc. Simply not all agents are in the position to act on their pro tanto responsibility. The second modification is to enable the agent to be *excused* from the responsibility altogether. This is relevant in the case of the brothels: it is not that those girls have a pro tanto obligation that is outweighed by other obligations of keeping themselves and their family safe. It seems more accurate to say that they do not have such a responsibility at all.

I thank Gunnemyr for these helpful amendments to the theory – but just as with accepting the general humanitarian principle in his first argument, these considerations do not diminish the normative significance of participation. Gunnemyr seems to throw out the baby with the bathwater. In fact, I think these two amendments only strengthen the idea of participation as the relevant criterion: note how the normative responses with respect to Williams' lorry driver were warranted, but modest. In the case of SI, we are only talking about *partial causal* responsibility. It would be inappropriate to let this then generate a heavy, inexcusable form of responsibility; Gunnemyr's amendments adjust the weight of the responsibility to the weight of the ground that generates it.

To sum: so far I have argued that participation in SI is the normatively relevant ground for a pro tanto and excusable remedial responsibility to combat that SI. The capacity to solve it could also be an appropriate ground, but this would be based on a general humanitarian principle, and such capacity is unlikely to even exist in the structural context.

3.3 Overdemandingness

My account so far might have given rise to two distinct worries of overdemandingness, that I briefly want to address here. The first concerns victims in general: even though we have allowed the responsibility to be pro tanto and excusable, is it *ever* sensible to identify a ground that makes victims responsible for changing the injustice they are the victims of? Is this not a form of victim blaming?

The latter is not the case, as we are talking about remedial responsibility and not blame (moral responsibility). And indeed, I do think such a position can be sensible. In Sandy's situation, for example, it makes sense to also ask of *her* to vote for certain policy makers, and to (say) support collective movements for the empowerment of renters. Perhaps we cannot ask as much from her as from a rich landlord, but she should not lean back while others take action, as she then too normalizes the problematic structures. Those victims who are not excused and whose pro tanto responsibility is not outweighed can still be asked to contribute to a solution. I even see two advantages of this position. First, quite pragmatically, the victims often have a privileged epistemic position: they can articulate what goes wrong and how they want to see it fixed. Secondly, more fundamentally, paternalism is avoided: only too often have situations been 'resolved' for victims of injustice without consulting them in this. It is in honor of their agency that their voice should be heard. Most instances of SI are precisely negations of the agency of the victims, and so any solution to it must start with acknowledging this agency.²⁵

A second worry of overdemandingness is that participation, as the relevant ground, is applicable to almost any agent in a large number of ways, manifesting itself in everyday actions like buying clothes or the driving of a car. This means that any agent is easily involved in a large number of SI's, being remedially responsible for all of those. But because we are finite beings with limited time and resources, we can never discharge all this responsibility. If we can never do what we are in principle required to do, is that not a flaw in the model?

I side with Young, who also discusses this worry, and argue against this. To start, the fact that we have responsibilities more than we can discharge also depends on *the amount of SI in the world*. If

²⁵ An extreme example of this is articulated by Franz Fanon, who argues that liberation from oppression can only be achieved by the victims themselves. Being granted freedom by the oppressor does not negate one's status of being oppressed (and lack of self-worth because of it) - it merely reaffirms it, as the freedom is *granted*. I do not endorse his claims here, but aim to show that my position is not too controversial. See *The Wretched of the Earth* (1961).

anything, *that* is the observation we should balk at, not the amount of responsibility it generates. “I suggest that in the contemporary world our political responsibility is limitless.” (McKeown, 2018, p.500) Secondly, as mentioned in section 2.3, Young formulates four “parameters of reasoning about responsibility” that determine *how much* one is responsible for: power (capacity), privilege, interest, and collective ability. She writes: “that we share responsibility [...] as contributors does not imply, however, that we should not distinguish *degrees* and kinds of responsibility”. (p.124, emphasis mine) What matters here is not Young’s specific choice of parameters, but the fact that distinguishing between degrees of responsibility, based on the circumstances of the agent, can ease the overdemandingness worry. If one scores lower on the chosen parameters, one has a lesser degree of responsibility.²⁶ And lastly, it is worth mentioning that “there is a conceptual difference between responsibility and duty that implies [...] openness. [With responsibility] It is up to agents [...] to decide what to do to discharge it within the limits of other moral considerations.” (Young, p.143) A duty prescribes a fixed action that needs to be undertaken; responsibility can be discharged according to the agent’s own discretion. I think these last two points together create enough breathing room for the overdemandingness worry to be softened.

3.4 Conclusion

In this chapter I have argued that the relevant criterion for establishing who, in fact, has remedial responsibility for structural injustice, is participation in the structures that cause the injustice. This participation has normative salience of its own, but not too strong: the responsibility it generates is pro tanto and excusable. The capacity to solve the problem could be a valid criterion too, but this hardly applies to the structural context, and would be based on a context-independent humanitarian principle. I have also argued that, in some cases, victims can also bear remedial responsibility, albeit often to a lesser degree. Allowing for degrees of responsibility, as well as leaving it to the discretion of the agent on how to discharge the responsibility, softens worries of overdemandingness.

²⁶ I would argue that if we allow the score on these parameters to reach zero, responsibility dissipates and we have caught Gunnemyr’s point that remedial responsibility can also be excused. However, Young seemingly did not allow for this: “*no one who participates* in processes that produce structural injustice *is exempt* from responsibility to join with others to change those structures.” (p.153, emphasis mine) I do not know why, given how this seems in line with the rest of her position.

Conclusion

In this thesis I have argued that participation in the social structures that cause structural injustice generates a pro tanto remedial responsibility to combat that injustice.

In spirit this conclusion, and this thesis in general, resembles Young's: she also used the concept of social structure to bring into view and scrutinize the background conditions of society. In doing this she identified all those agents connected to the structural injustice as responsible. She firmly rejected blameworthiness for it, while adding nuance by distinguishing between different degrees of responsibility.

But to this I have added new elements, trying to sharpen some ideas and distinctions. I described social structure more generally, employing Sangiovanni's account of abstract schemas and resource distributions mutually shaping each other. I used this to make explicit the causal link between social structure and injustice, using Hart and Honoré's theory of causality. In the second chapter I adopted Miller's distinction between causal, outcome, moral and remedial responsibility. This distinction served us well: applying the structural context as it was constructed in the first chapter, we could establish partial causal responsibility and reject outcome and moral responsibility. This partial causal responsibility was then identified in the third chapter as the relevant ground for remedial responsibility, using Williams' thought experiment of the unfortunate lorry driver. I also diverged from Young in other respects: I followed Goodin and Barry in conceiving of taking remedial responsibility as 'taking a moral stand', as opposed to a more consequentialist view. I allowed blame for agents who do not even try to shoulder their responsibility, and I adopted Gunnemyr's suggestion that remedial responsibility for structural injustice should be pro tanto and excusable. All in all I have tried to assemble useful ideas from many different thinkers into one improved, coherent whole.

Throughout the thesis I left aside some interesting, related topics due to space constraints. I want to briefly mention three of these here, aiming to sketch some possible directions for future research. There is the question of which specific actions the remedial responsibility for structural injustice prescribes (perhaps combined with an analysis of collective actions; see Schwenkenbecher, 2021). There is the notion of social structure, the analysis of which could be connected to related ideas in continental philosophy. And there is the discussion on epistemic conditions of responsibility: (how much) does it matter whether one knows about one's responsibility? Answers to these inquiries could substantially strengthen the structural account.

Acknowledgements

I want to briefly thank my supervisor Wouter Kalf, who always made me leave his office with more confidence than when I came in. I also thank Catelijn van Berkel for enduring and backing me throughout the writing process. I enjoyed the support from many friends who engaged me in discussions on this topic, as well as a few trusted peer-reviewers: Maarten van Tunen, Tim van der Meulen, Oran Dan, Cees Timmers and Vincent Rouwmaat.

Bibliography

- Abdel-Nour, F. (2018). Responsibility for structural injustice. *Ethics & Global Politics*, 11(1), 13–21.
- Aragon, C., & Jaggar, A. M. (2018). Agency, Complicity, and the Responsibility to Resist Structural Injustice. *Journal of Social Philosophy*, 49(3), 439–460.
- Ayala-López, S. (2018). A Structural Explanation of Injustice in Conversations: It's about Norms. *Pacific Philosophical Quarterly*, 99(4), 726–748.
- Barry, C., & Ferracioli, L. (2013). Young on Responsibility and Structural Injustice. *Criminal Justice Ethics*, 32(3), 247–257.
- Barry, C., & Macdonald, K. (2016). How should we conceive of individual consumer responsibility to address labour injustices? In F. Milman-Sivan, H. Lerner, & Y. Dahan (Eds.), *Global Justice and International Labour Rights* (pp. 92–118). Cambridge University Press.
- Beck, V. (2020). Two forms of responsibility: Reassessing Young on structural injustice. *Critical Review of International Social and Political Philosophy*, 1–24.
- Chung, R. (2021). Structural health vulnerability: Health inequalities, structural and epistemic injustice. *Journal of Social Philosophy*, 52(2), 201–216.
- Collins, S. (2018). When does 'Can' imply 'Ought'? *International Journal of Philosophical Studies*, 26(3), 354–375.
- Digeser, P. E. (2018). Motivation and reconciliation in Catherine Lu's conception of global justice. *Ethics & Global Politics*, 11(1), 6–12.
- Eckersley, R. (2016). Responsibility for Climate Change as a Structural Injustice. In T. Gabrielson, C. Hall, J. M. Meyer, & D. Schlosberg (Eds.), *The Oxford Handbook of Environmental Political Theory* (pp. 346–361). Oxford University Press.
- Goldman, A. I. (1999). Why Citizens Should Vote: A Causal Responsibility Approach*. *Social Philosophy and Policy*, 16(2), 201–217.
- Goodin, R. E., & Barry, C. (2021). Responsibility for structural injustice: A third thought. *Politics, Philosophy & Economics*, 20(4), 339–356.
- Gunnemyr, M. (2020). Why the Social Connection Model Fails: Participation is Neither Necessary nor Sufficient for Political Responsibility. *Hypatia*, 35(4), 567–586.
- Hart, H. L. A., & Gardner, J. (2008). *Punishment and Responsibility: Essays in the Philosophy of Law* (2nd ed.). Oxford University Press.
- Hart, H. L. A., & Honoré, T. (1959). *Causation in the Law*. Oxford University Press.

- Haslanger, S. (2016). What is a (social) structural explanation? *Philosophical Studies*, 173(1), 113–130.
- Kutz, C. (2000). *Complicity: Ethics and Law for a Collective Age*. Cambridge University Press.
- Lepora, C., & Goodin, R. E. (2013). Complicity and its Conceptual Cousins. In C. Lepora & R. E. Goodin, *On Complicity and Compromise* (pp. 31–58). Oxford University Press.
- Lu, C. (2017). *Justice and Reconciliation in World Politics*. Cambridge University Press.
- (2018). Responsibility, Structural Injustice, and Structural Transformation. *Ethics & Global Politics*, 11(1), 42–57.
- McKeown, M. (2015). *Responsibility without Guilt: A Youngian Approach to Responsibility for Global Injustice* [Doctoral]. UCL (University College London).
- (2018). Iris Marion Young’s “Social Connection Model” of Responsibility: Clarifying the Meaning of Connection. *Journal of Social Philosophy*, 49(3), 484–502.
- (2021). Structural injustice. *Philosophy Compass*, 16(7).
- Miller, D. (2001). Distributing Responsibilities. *Journal of Political Philosophy*, 9(4), 453–471.
- (2007). *National Responsibility and Global Justice*. Oxford University Press.
- Moore, M. (2019). Causation in the Law. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2019). Metaphysics Research Lab, Stanford University.
<https://plato.stanford.edu/archives/win2019/entries/causation-law/>
- Nussbaum, M. (2011). Foreword. In I. M. Young, *Responsibility for Justice*. Oxford University Press.
- Parekh, S. (2012). Does ordinary injustice make extraordinary injustice possible? Gender, structural injustice, and the ethics of refugee determination. *Journal of Global Ethics*, 8(2–3), 269–281.
- Parfit, D. (1986). *Reasons and Persons*. Oxford University Press.
- Sangiovanni, A. (2018). Structural Injustice and Individual Responsibility: Structural Injustice and Individual Responsibility. *Journal of Social Philosophy*, 49(3), 461–483.
- Sankaran, K. (2021). “Structural Injustice” as an analytical tool. *Philosophy Compass*, 16(10).
- Schwenkenbecher, A. (2021). Structural Injustice and Massively Shared Obligations. *Journal of Applied Philosophy*, 38(1), 23–39.
- Tognazzini, N., & Coates, D. J. (2021). Blame. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2021). Metaphysics Research Lab, Stanford University.
<https://plato.stanford.edu/archives/sum2021/entries/blame/>
- Williams, B. (1981). *Moral Luck: Philosophical Papers 1973–1980*. Cambridge University Press.
- Young, I. M. (2011). *Responsibility for Justice*. Oxford University Press.

Zheng, R. (2018). What is My Role in Changing the System? A New Model of Responsibility for Structural Injustice. *Ethical Theory and Moral Practice*, 21(4), 869-885.

----. (2019). What Kind of Responsibility Do We Have for Fighting Injustice? A Moral-Theoretic Perspective on the Social Connections Model. *Critical Horizons*, 20(2), 109-126.

----. (2021). Moral Criticism and Structural Injustice. *Mind*, 130(518), 503-535.