

An Account of Structural Change

Mauro Rossi (UQAM) and Stuart Yasgur (Economic Architecture)

Draft: 13/09/2026

1. Introduction

In 2010, the U.S. Supreme Court's decision in *Citizens United v. Federal Election Commission* removed limits on corporate spending in elections. What happened next was entirely predictable: money flooded into political campaigns¹ and corporations gained outsized influence over policy, undermining the conditions for meaningful democratic representation. This was no accident, but the result of a precise logic: politicians need money to win elections, wealthy donors have money and interests to protect, so each side found it rational to engage with the other. No amount of moral exhortation, civic pressure or legal challenge has managed to change things. In fact, the more entrenched this dynamic became, the harder it became to reverse it, as those who benefit most from the new legislation are precisely those with the resources and influence to resist its reform.

Citizens United is a paradigmatic instance of structural change:² it transformed the structure of the U.S. campaign finance system. That the change was structural is partly what makes it so significant, and so difficult to reverse. Social structures profoundly influence and constrain individuals' choices, steering systems toward certain outcomes and generating a sense that those outcomes are, to some degree, beyond individual control. This is precisely what *Citizens United* produced: a political environment in which the rules of the game systematically favour those with the most resources, and in which ordinary citizens feel largely powerless and distrustful.

Structural change is not always negative. The notion is frequently invoked in academic and policy discussions as a remedy to structural problems. To address climate change or structural racism, it is said, we need structural change. The idea is intuitive: if structural problems are rooted in and sustained by particular social structures, addressing them requires changing those structures. In this sense, structural change is often precisely what is needed and what we should aim for. In its positive as in its negative guise, what makes structural change significant is the degree to which it shapes outcomes and is resilient to change. It transforms things in ways that are more far-reaching and durable than other kinds of change.

¹ The figures are striking. *OpenSecrets* reports: "Nonprofit groups funded by secret donors disclosed less than \$1 million in spending to the FEC during the first half of the two-year election cycle prior to the Supreme Court's *Citizens United* in January 2010. That spending increased by more than one-hundredfold in a matter of months, topping \$138 million before Election Day." (See <https://www.opensecrets.org/news/2020/01/dark-money-10years-citizens-united/>.) By 2012, that figure had reached \$313 million. (See <https://www.opensecrets.org/news/reports/a-decade-under-citizens-united>.) By 2024, dark money spending on federal elections was estimated at \$2 billion, double the total spent in 2020. (See <https://www.brennancenter.org/our-work/analysis-opinion/new-study-shows-runaway-influence-dark-money-politics>.)

² Throughout the paper, we focus on *social* structural change, that is, change pertaining to social structures. We use the term 'social' structure broadly, encompassing 'economic' and 'political' structures. For simplicity, in what follows we will often omit this qualification.

Despite its importance, the notion of structural change is rarely given precise content. Little attention, in particular, has been paid to the question of what distinguishes structural from non-structural change.³ This is a surprising lacuna, given that the question clearly matters for two distinct reasons. First, the concept of structural change is meant to pick out changes that alter structures in fundamental ways. Without a clear account of what makes a change structural, we cannot reliably identify such changes or explain why they have the effects they do. Second, if structural change is what structural problems require, we need to know what it takes to produce a change that is genuinely structural. An adequate account must therefore be descriptively and explanatorily adequate as well as action-guiding, articulating the conditions under which a change is structural and under which it is not. The goal of this paper is to offer such an account

This task immediately runs into a methodological difficulty. Generally speaking, a structural change is a change *of* the structure, one that produces a qualitatively different structure, rather than a change that occurs *within* an otherwise unaltered structure, such as a parametric adjustment of some of its elements. From this, one may infer that to develop an account of structural change, one must first have a *specific* account of social structures. The problem with this approach is that it makes the resulting account of structural change hostage to the fortunes of the underlying theory of social structures. For this reason, we propose a different methodology. Our approach combines reflection on *general* and widely accepted features of social structures – features that any plausible account must accommodate and explain – with a systematic analysis of concrete case studies. This is, in effect, the method of reflective equilibrium applied to the problem of structural change.

Accordingly, Section 2 presents a series of real-world cases, some clear instances of structural change, others clearly non-structural, and another with a more uncertain status. A satisfactory account must explain and vindicate our considered judgments about the paradigmatic cases as well as provide a framework for making a reflective judgment about the ambiguous one. Section 3 develops our theoretical account. We build from three widely recognized features of social structures: they determine the possibility space for agency (Haslanger 2016), they make certain outcomes highly likely if not nearly unavoidable, and they exhibit significant stability. We provide an interpretation of each of these features and explain what is required to modify them to produce a structural change. We argue that a change counts as structural if and only if it satisfies three conditions: (1) it alters the set of possible strategies that the relevant actors can choose; (2) it ‘programs’ (in Jackson and Pettit (1992)’s sense) for a new outcome, or for an outcome that was previously unprogrammed; and (3) it generates feedback loops that make the change self-sustaining.

Throughout the paper, we assume three things. First, a change – whether structural or non-structural – occurs through the introduction of one or more factors. We use the term ‘introduction’ liberally, to indicate either the addition of a new factor or the subtraction of a previously existing factor. Whenever a structural change occurs *by virtue of* the introduction of a particular factor, we call that

³ To date, only one unpublished manuscript, by Sally Haslanger, has been devoted to this issue. This is all the more surprising given that, at least since the posthumous publication of Iris Marion Young’s *Responsibility for Justice* (2011), there has been a renewed interest in the related topic of structural injustice. Debates in this area have, however, centered on two different issues: the scope of individual responsibility for structural injustices (see Young 2011; Sangiovanni 2018; Zheng 2018; McKeown 2024), and whether structural injustice constitutes a genuine moral wrong or a mere ‘bad’ (see Sankaran 2021; Estlund 2024; Sankaran & Monaghan 2025; Collins 2025; and Estlund 2026).

factor a ‘structural factor’.⁴ Such a factor performs a ‘structuring’ function: its introduction results in, and is the ground of, a qualitatively different structure. Second, we are pluralistic about structural factors. We think that a structural change may occur by virtue of the introduction of laws, policies, and institutions, but also of technologies, programs, organizations, and possibly other things. Third, we assume that structures can be located at different levels, or, equivalently, that there exist different structural levels. Some structures are more ‘local’ (e.g., the structure of a city council), others more ‘global’ (e.g., the political structure of a country). In addition, some structures are embedded in other structures (e.g., the structure of the labour market is embedded in the overall structure of a country’s economy). These distinctions are important when it comes to structural change, for they imply that a relevant factor may be structural at one level, but not at another. For example, the introduction of new labour laws may alter the structure of the labour market, but not the overall structure of the economy. A given structural factor is therefore ‘structural’ only relative to a specific structure. That said, we maintain that the *conditions* determining whether a change is structural are invariant across levels. What varies is simply whether those conditions are satisfied in relation to a given structure. The aim of this paper is to spell out precisely what those conditions are.

2. Case studies

Here are the examples we consider, in addition to *Citizens United*.⁵

European Single Market. The European Single Market encompasses the 27 member states of the European Union as well as four non-EU countries that participate through the European Economic Area. It was officially launched in 1993. The Single Market guarantees the free movement of goods, services, capital, and people across borders, with the goals of stimulating economic growth, enhancing the EU’s global competitiveness, and deepening political and social integration. The core idea is that businesses benefit from access to a larger, integrated market in which they can scale more easily and operate with fewer regulatory barriers, while consumers gain from a wider selection of goods and services, often at lower prices due to increased competition. Citizens are also free to live, work, and study in any participating country, fostering both economic mobility and cross-cultural exchange. Since 2002, most EU countries have adopted a common currency, the euro, which further reduces transaction costs and eliminates exchange rate risks, reinforcing the economic cohesion of the Single Market.

Federal Reserve. The Federal Reserve was established in 1913 to address recurring banking crises, particularly the financial panics that frequently destabilized the U.S. economy. Before its formation, the collapse of one bank often triggered widespread panic, prompting depositors at otherwise stable banks to withdraw their funds, thus exacerbating financial instability. To mitigate these crises, banks needed access to a centralized source of emergency liquidity. Private clearinghouses sometimes coordinated responses to stabilize the banking system, but their efforts were often reactive, inconsistent, and limited in scope. The Federal Reserve was designed to serve

⁴ The italics is important to distinguish genuine structural factors from factors that merely *trigger* a structural change (i.e., that put in motion a series of events that eventually lead to a structural change) but that are not the *bases* of such a change.

⁵ One caveat: The examples listed in this section are about changes and innovations that are more complex than what our presentation suggests. For reasons of space, here we can only provide a very succinct and simplified presentation.

as a lender of last resort, providing liquidity during financial crises to prevent bank failures and broader economic collapse. Over time, it took on a more active role in managing economic stability, primarily by influencing interest rates through open market operations, the federal funds rate, and other monetary policy tools. By raising rates, the Fed can curb excessive borrowing and spending to prevent inflation, while lowering rates encourages investment and consumer spending to stimulate economic growth during downturns.

Netflix. Netflix is a subscription-based streaming service. In January 2022, it raised the price of its standard subscription plan in the U.S. from \$13.99 to \$15.49 per month, and its premium plan from \$17.99 to \$19.99. The decision was driven by rising content costs and slowing subscriber growth. While some subscribers cancelled their accounts in the wake of Netflix's decisions, overall revenue grew. The change led competitors such as Disney+ and HBO Max to reconsider their own pricing strategies, and several followed with increases of their own in the subsequent months. Netflix's pricing model has remained broadly stable since.

Voluntary Carbon Offsetting. In the early 2000s, several organisations started to offer individuals the possibility of purchasing carbon offsets to compensate for their personal greenhouse gas emissions, most notably, those generated by air travel, car use, and household energy consumption. The idea behind these schemes was simple: consumers paid a fee that funded projects designed to reduce or capture an equivalent quantity of carbon dioxide elsewhere, for instance, through reforestation initiatives or renewable energy installations. By the mid-2020s many major airlines had integrated carbon offset options directly into their booking processes. Participation over the years has remained modest relative to the scale of the emissions involved, but the schemes have attracted a stable base of users and generated significant revenue for offset project developers.

Québec Bill 83. The province of Québec, Canada, has faced a persistent shortage of both family doctors and specialists in recent years. This issue is particularly severe in rural regions, which are often remote and sparsely populated. Many factors contribute to this situation. Access to medical school in Québec is restricted by a quota system, and while the cost of training a doctor is higher than in other Canadian provinces, tuition fees remain significantly lower.⁶ In this context, many doctors choose to study in Québec but then leave for other jurisdictions, attracted by better opportunities.⁷ Furthermore, Québec's healthcare system is a hybrid of public and private services, and since 2020, the number of doctors working exclusively in the private sector has risen by 70%, a trend particularly pronounced among new doctors.⁸ One frequently cited cause is the government's annual plan, which sets target numbers for physician positions within each administrative region and effectively regulates where new doctors can practice, often directing them to understaffed areas where they face exhausting working conditions. Given these constraints, many new doctors choose to leave the public system for the private sector, where they have the freedom to choose their location, along with better working conditions and higher pay.⁹ In response to this situation, the Québec government has considered a bill that would require all doctors trained

⁶ For example, domestic students at the University of Toronto pay approximately C\$23,000 per year, whereas tuition at McGill University in Montréal is around C\$3,000. See <https://globalnews.ca/news/10899829/que-doctors/>.

⁷ In 2024, for instance, 2,355 Québec-trained doctors were practicing in Ontario, including 1,675 who attended McGill University. See <https://globalnews.ca/news/10899829/que-doctors/>. For a dissenting view about the facts on the ground, see <https://ici.radio-canada.ca/nouvelle/2138691/exode-medecins-quebec-canada-dube>.

⁸ See <https://globalnews.ca/news/10899829/que-doctors/>.

⁹ See <https://www.cbc.ca/news/canada/montreal/record-number-quebec-doctors-leave-public-sector-1.7273350>.

in the province – both general practitioners and specialists – to work in the public healthcare system for the first five years of their careers. Those who fail to comply could face fines of up to C\$200,000 *per day*.¹⁰

We take the first two cases to be paradigmatic instances of structural change, like *Citizens United*. Indeed, if these did not count as structural, it would be difficult to imagine what would. The creation of the European Single Market fundamentally transformed the organization of the European economy. Likewise, the establishment of the Federal Reserve reshaped the structure of the U.S. financial system. Thus, the Single Market, the Federal Reserve, as well as the permissive campaign finance regime that resulted from *Citizens United*, all qualify as structural factors in our sense.

We think instead that the next two cases are paradigmatic instances of *non*-structural change. The relevant factor in the *Netflix* example is a higher subscription price. The structure of interest is the digital media and entertainment market. The price increase did not introduce a new element into that structure; it simply modified the value of an existing one, with the aim of influencing consumer behaviour. This amounts to a parametric change. It is therefore best understood as a non-structural change. The *Voluntary Carbon Offsetting* case is more interesting. The possibility of purchasing carbon offsets was a genuine new element introduced into the system. Yet, it did not change the market for carbon-intensive goods and services at a fundamental level. The scheme simply attempted to induce consumers to bear part of the environmental costs of their consumption habits, without altering the structures that make fossil fuels the dominant and default source of energy supply. This too, therefore, represents a non-structural change.

Québec Bill 83 is more ambiguous. It is difficult to have a firm pre-theoretical opinion about it, which is precisely what makes it a valuable test case.

It is worth emphasizing that these should be understood as ‘considered judgments’, i.e., judgments that, upon reflection, should be shared by competent observers with an intuitive grasp of what counts as the relevant structure in each case. They are also meant to be compatible with different theories of social structures, at least insofar as those theories recognize that structures exist at multiple levels and are pluralistic about structural factors, e.g., they do not restrict structural change to changes in laws or policies. Crucially, these judgments do not presuppose a full-fledged theory of structural change; rather, they provide a starting point for developing one. We begin to do that in the next section.

3. Our Account of Structural Change

3.1. Changing the available strategies – a.k.a. the strategy set condition

There is broad consensus in the literature that social structures determine what Sally Haslanger has dubbed the “possibility space for agency” (2016: 127), that is, the possibilities open to the agents

¹⁰ See <https://www.cbc.ca/news/canada/montreal/quebec-bill-doctors-health-care-1.7400228>.

who operate within those structures.¹¹ This immediately suggests a first condition for structural change: for a change to count as structural, it must alter the “possibility space for agency”.

We agree. The crucial question, however, is how to understand the “possibility space for agency”. It is not always clear or explicit in the literature what this expression refers to. Most authors appear to interpret the relevant possibilities in terms of the *options available* to the individuals within the structure.¹² On this view, a structural change requires a change in the options available to the relevant agents.

We think this is on the right track but not quite correct. We propose two modifications. First, the “possibility space for agency” should be defined over *strategies* rather than options. Social structures do not merely determine single and isolated options; rather, they shape the agents’ broader plans of action and the objectives they can pursue. Second, what matters for structural change is a change in the agents’ *possible* strategies, not in their *available* ones. The former are the strategies that the structure itself makes possible; the latter are the subset of possible strategies that remain feasible at a given point in time within the structure. The first condition for structural change therefore becomes: for a change to count as structural, it must alter the set of possible strategies open to the relevant agents.

An analogy with chess may help clarify these distinctions and their significance. Chess is a two-player strategic board game in which each player controls sixteen pieces – king, queen, rooks, bishops, knights, and pawns – each with its own movement rules. The objective is to place the opponent’s king in checkmate. The game unfolds through a sequence of moves by which players enact their strategies. A *move* is a single concrete action at a specific moment, such as advancing a pawn. Moves should be distinguished from options: an *option* is an ordered pair comprising a move and the outcome it produces. This explains why one and the same move need not constitute the same option in different situations: moving a pawn may result in capturing a knight in one case and defending the queen in another. A *strategy*, by contrast, is a broader plan guiding a player’s choices over the course of the game, in view of the objective of winning the game. *Possible* strategies are all those that can be played in accordance with the rules of chess. *Available* strategies are those that remain feasible for a player at a particular point in the game, given what has already occurred.

When a player makes a move, this typically changes the options and strategies available to both players. However, it does not alter what is possible to do according to the rules. Castling kingside, for instance, is part of a possible strategy since the rules permit it, but it may no longer be part of a strategy available to a player who has already moved their king or rook. A move that alters the available options and strategies is clearly a change *within* the game, not *of* the game itself. So, what does it take to change chess ‘structurally’, that is, to turn it into a different kind of game?

One way is to change the permissible moves, e.g., allowing bishops to move like knights, or pawns to move backward. But this is not the only way. Chess can be turned into a different game without altering any of the permissible moves. For example, a change of this sort can be achieved by changing the number of players, the shape of the board, or what it takes to win the game. The

¹¹ In addition to Haslanger, see, for example, Young 2011; McKeown 2021; Ross 2024; and Hindriks 2025.

¹² See Young 2011, p. 53-55; Haslanger 2016, pp. 116-117, pp. 123-125, p. 128. Other authors interpret the “possibility space for agency” in terms of “actions” (see McKeown 2021, p. 3) or “outcomes” (see Ross 2024, p. 173).

crucial thing to notice is that all such changes alter the set of possible strategies open to players: the broader plans of action they can pursue and the objectives they can aim at. This suggests that the “possibility space” in chess is, most fundamentally, the space of possible strategies, and a structural change in chess requires a change in that space.

The same point applies to social structures. We acknowledge that the analogy is imperfect: in social life there may be no explicit or formal rules establishing what is possible for agents to do. But it is close enough to be illuminating. What fixes the space of possible strategies in social life is the ensemble of what we have called structural factors, that is, the factors that perform ‘structuring’ work, such as institutions, laws, policies, technologies, etc. Changing a social structure requires changing that space. A change that leaves it intact, however significant its other effects, is a change within the structure, not of it.

We can see that this condition is satisfied in the paradigmatic cases of structural change presented above. Consider *Citizens United*. The U.S. Supreme Court’s decision clearly changed the actions that the relevant agents are permitted to do. In particular, it allowed corporations to spend unlimited funds on elections. More fundamentally, however, it altered the strategies open to the relevant agents. Corporations can now reallocate resources away from traditional lobbying toward direct electoral campaigning. Political candidates and parties, in turn, can seek potentially unlimited funds from a limited pool of wealthy donors, rather than having to diversify their funding sources. The decision has thereby transformed the strategic landscape of electoral competition.

A similar transformation occurred with the creation of the European Single Market. This reform expanded the strategy set of the relevant economic actors. Firms gained access to new possibilities for organising production and distribution across member states, reconfiguring supply chains and competing for a wider consumer base. Workers and students can now seek employment or study opportunities in other EU countries under the free movement provisions. The change thus expanded the set of possible strategies of the agents within the European economic system.

The first condition also explains why *Netflix* is a *non*-structural change. The increase in subscription prices did not modify the set of possible strategies open to the relevant agents. When Netflix raised its prices, it simply opted for one possible strategy open to them within the digital media market. This immediately altered the strategies that remained available to other agents: budget-constrained consumers faced potential trade-offs in their consumption strategies, and rival firms had to decide whether to engage in price competition or pursue different strategies. But this is precisely analogous to a move in chess that restricts the strategies available to the other player without altering their possible strategies. It is a change *within* the structure, not *of* the structure, and as such is a non-structural change.^{13,14}

¹³ These considerations explain why parametric changes are (typically) non-structural – a point that underlies our intuitive case for considering *Netflix* a non-structural change. A parametric variable is one that can take various values within a given range. That range is one of the elements that determine the set of *possible* strategies open to the relevant agents. A subsequent variation in the value of the parameter, e.g., a price increase, merely alters the strategies that remain *available* to agents given that variation; it does not expand or transform the space of possible strategies defined by the range itself. The first condition thus neatly accounts for our considered judgment that *Netflix* is a non-structural change: a price adjustment modifies a parameter within an existing structure rather than altering the structure itself.

¹⁴ This also suggests that it is not enough for a measure to *merely* change agents’ incentives for it to bring about a structural change.

3.2. ‘Ensuring’ the outcomes – a.k.a. the robustness condition

In the previous subsection, we argued that a necessary condition for a change to count as structural is that it modifies the set of possible strategies that the agents can choose.¹⁵ This is not sufficient. Adding an insignificant new soda to the already overcrowded soda market may alter the agents’ possible consumption strategies, yet, intuitively, this change does not count as a structural change of the soda industry. It seems that, to count as structural, a change must also have *some* impact on the decisions the agents actually make, not just on their possible strategies. But which kind of impact? It is tempting to say that for a change to count as structural, it must bring about a change in *outcome*. But this condition is too broad. As we will argue below, there seem to be cases of structural change where a change in the set of possible strategies affects only *how* an outcome is realized, rather than *which* outcome is realized.

To make progress, let us consider another important feature of social structures. Many authors have argued that, beyond shaping the options available to agents, social structures can make certain choices and their corresponding outcomes so likely as to appear almost inevitable. It is as if social structures *compel* particular decisions and generate specific outcomes with near ineluctable certainty. It would be too strong to say that social structures ‘determine’ those choices and outcomes – a claim few contemporary theorists are willing to endorse. Rather, their influence is more varied. Lauren Ross (2024, pp. 173-174) has recently suggested that social structures operate as “causal constraints”. In some circumstances, the constraint is so “extreme” that it is essentially impossible for an agent to exhibit any other choice behaviour than their actual behaviour. In other circumstances, however, social structures operate simply as “strong constraints”, in the sense that they make certain choices more likely or more rational than others. These remarks are representative of a common position in the literature and point to an additional condition on structural change. Formulated in terms of strategies: for a change to be structural, it must alter the strategies that are either more likely or more rational for individuals to choose.

As stated, however, this proposal will not do. If Jack knows that Mary has idiosyncratic psychological vulnerabilities, he can change her strategy set in a way that makes a particular strategy more likely or more rational for her. Yet, this does not necessarily constitute a structural change, since the effect of Jack’s intervention depends entirely on Mary’s specific psychological profile and would not extend to other individuals in relevantly similar situations. This suggests a more general point. Social structures do not merely make certain strategies more likely or rational *given a specific set of circumstances*. Rather, they exhibit a kind of counterfactual robustness: they would still make those strategies more likely or rational *even if the circumstances were somewhat different*, such as if the individuals involved had partly different psychological dispositions. This idea is frequently mentioned in passing in the literature,¹⁶ but we need a more precise theoretical statement to capture it.

¹⁵ There are various ways to change the possible strategies open to the agents. One approach, which we have not explored in the previous sub-section, consists in *removing* some types of agents from the system. Consider, as an especially worrisome example, Trump’s plan to shut down the Consumer Financial Protection Bureau, which is likely to lead to significant changes in the strategies open to consumers and financial institutions.

¹⁶ Haslanger (2016), for instance, refers to this point by saying that structural explanations are more stable than non-structural explanations: the former do not depend on contingent features of the event (e.g., the particular ball that was inserted in the treat falling downhill), but would hold even if these features had been different (e.g., even if the ball in

We propose to find inspiration in Frank Jackson and Philip Pettit's (1992) account of social structural explanations. Jackson and Pettit's starting point is a puzzle that arises out of three assumptions: the first is that social explanations are explanations of social facts; the second is that social facts ultimately obtain in virtue of the intentional attitudes of the individuals involved; the third is that social *structural* explanations are macro-level explanations of social facts that do not make reference to the individuals' attitudes. The puzzle they try to solve is how such non-psychological explanations can be useful if social facts are ultimately caused by psychological facts. The key, they think, lies in understanding how structural explanations work.

Jackson and Pettit begin by distinguishing between process explanations and program explanations. Process explanations account for events by identifying their actual causes. For example, a process explanation of why a flask cracked would cite the collision of a particular group of molecules with a molecular bond on the flask's surface at a velocity sufficient to break it. By contrast, program explanations identify a condition whose obtainment 'programs' for the occurrence of these events, where to program for an event is to ensure, or make it more likely, that there will be *some* causes to produce the event in question; if not the actual cause, another one. For instance, a program explanation of why the flask cracked might cite the fact that the water inside reached boiling point, making it likely that molecular collisions sufficient to cause cracking would eventually occur – if not the particular collision that did, then another one.

Jackson and Pettit argue that social structural explanations are a kind of program explanations, characterized by having social facts as their explananda. Given that social facts ultimately obtain in virtue of the intentional attitudes of the individuals involved, social structural explanations take the following form: they explain social facts by identifying a factor that ensures, or makes it more likely, that there will be *some* combination of intentional attitudes at the individual level sufficient to cause the social fact in question; if not the actual combination, a different one. For example, consider a structural explanation of rising crime rates that attributes the increase to rising unemployment. This explanation holds that unemployment ensures, or increases the likelihood, that *some* psychological mechanisms – such as anger at one's condition, a desire to escape poverty, or a need to provide for one's family – will lead individuals to commit crimes. Thus, even if the specific motivations vary among individuals, unemployment 'programs' for the emergence of crime.

These considerations allow Jackson and Pettit to explain why social structural explanations are useful: they provide information that process explanations do not provide. First, they show how certain events are not merely the product of contingent psychological factors but have a kind of counterfactual robustness: they would have occurred even if the individuals' psychological make-up had been somewhat different. Second, by identifying particular 'programming' factors, they explain what makes the occurrence of these events robust, offering deeper insight into the conditions under which they arise.

Let us go back to social structures. We saw that social structures seem to make certain choices and outcomes almost inevitable, and this, regardless of the specific individuals that are involved. We

the treat had been different). We prefer to describe these explanations as 'robust', reserving the term 'stable' for a distinct feature of social structures.

propose to capture this idea by saying that social structures ‘program’ for certain outcomes in Jackson and Pettit’s sense, meaning that they involve factors that ensure or make it more likely that there will be *some* combination of intentional attitudes and behaviors at the individual level sufficient to cause the outcome in question; if not the actual combination, a different one.¹⁷

We can make this idea more precise. Social structures program for particular outcomes by shaping the individuals’ choices within their possibility space for agency, that is, the domain of possible strategies open to them. More specifically, social structures ensure or make it more likely that, in one way or another, individuals will choose some of these possible strategies rather than others, thus producing a particular outcome. *Typically*, we add, they do so by making some of these strategies more *rational* than others for the individuals involved, and this, regardless of those individuals’ identities, histories, or actual motivations.

This yields a second condition for structural change: for a change to count as structural, it must have ‘programming’ features: it must ensure or make it more likely that certain outcomes will robustly occur as a result of the individuals’ choices – typically, by making some strategies more rational across a wide range of individual variations.

The *Citizens United* case illustrates this clearly. The Supreme Court’s decision did not merely open new strategies for the relevant actors; it made those strategies rational for them in a counterfactually robust way. It is now instrumentally rational for any corporation with sufficient resources to inject large sums of money into elections – whether it is an oil company or an environmental organisation, whether it is motivated by ideology or business interests. By funding politicians, corporations can indeed ensure that their interests will be adequately protected. Likewise, soliciting contributions from wealthy donors is now a rational funding strategy for any politician seeking election – whether Republican or Democrat, whether in politics to enrich themselves or to promote social justice. Money is a strong predictor of electoral success. Courting large donors is thus a rational response regardless of who the politician happens to be. By making these strategies rational in such a robust way, *Citizens United* has thus effectively ‘programmed’ for a specific outcome: one where corporations and special interest groups exercise disproportionate influence over elections and policy decisions.

Similar considerations apply to the *Federal Reserve* case. The creation of the Federal Reserve introduced a new kind of agent into the financial system, capable of pursuing different possible strategies, such as maintaining low interest rates to stimulate growth or adopting a more inflation-cautious approach. This, in turn, created new strategic opportunities for banks, policymakers, consumers, and firms, all of whom now had to calibrate their own strategies in response to the Federal Reserve’s choices. But the creation of the Federal Reserve also had programming effects. By adopting particular interest-rate policies, the Federal Reserve ensures that certain combinations of strategies become rational for the relevant agents, regardless of their individual histories, motivations, or personal characteristics. The Fed is thus able to steer the system toward particular outcomes in a way that is robust across a wide range of individual variation. In this way, the creation of the Federal Reserve has ‘programmed’ for greater stability in the financial system.

¹⁷ We do not deny, however, that the programming factors also play a (higher order) *causal* role in producing the outcome.

We can test our second condition in the other direction by examining whether it fails to hold in paradigmatic cases of non-structural change. Consider *Voluntary Carbon Offsetting*. The possibility of purchasing carbon offsets altered consumers' possible consumption strategies. It also contributed to more desirable environmental outcomes. But these schemes have neither ensured nor made it more likely that purchasing offsets constitutes a *robustly* rational strategy. At most, it is rational for consumers whose choices are especially sensitive to environmental considerations – and even then, only under certain conditions, e.g., when the additional cost is not excessive. The schemes therefore lack the programming features characteristic of structural factors: specifically, they do not ensure or make it more likely that carbon emissions will be systematically reduced or compensated regardless of the identity and motivations of the relevant agents.

We conclude by noting that the second condition can also explain how there can be structural changes without changes in outcome, a possibility that we noted at the start of this subsection. This can happen when the introduction of a new factor makes the occurrence of an outcome that was previously the product of merely contingent circumstances robust, ensuring its occurrence across varying circumstances. Consider the following example. The emergence of food deserts is often regarded as a *structural* barrier to healthy food consumption. Imagine a scenario where, initially, residents of a low-income neighbourhood have only a limited preference for healthy food. At this stage, their dietary choices are shaped by their tastes, but they still have the option to seek out healthier food locally. The resulting outcome – residents do not consume healthy food – is thus contingent on their current tastes. Once healthy food retailers leave the area, however, individuals with limited financial resources face prohibitively high costs, whether in time or money, to access nutritious food. Consumption strategies involving the purchase of healthy food cease to be possible within the neighbourhood's food market. Even if residents' preferences change and they come to desire healthier food, their only possible choice is to purchase from nearby fast-food outlets. The outcome remains the same – residents do not consume healthy food – but its status has changed: whereas its occurrence was previously contingent on individual preferences, it is now robust and independent of those preferences. Thus, the departure of healthy food retailers has effectively *programmed* for an outcome that was previously unprogrammed. Note that the departure of healthy food retailers is a possible 'move' within the broader food market, which does not alter *that* market's structure. It does, however, have an important effect on the structure of the low-income neighbourhood's food market: it alters the consumption strategies that are possible for agents operating within that market and thereby programs for a particular outcome. Thus, if the present case qualifies as a structural change of that local market, it does so, in part, because of this programming feature.

3.3. Changing the feedback loops – a.k.a. the stability condition

So far, we have argued that a structural change must satisfy two conditions: it must alter the set of possible strategies open to the relevant agents, and it must program for particular outcomes. These conditions are individually necessary but not jointly sufficient. To see why, compare the following three scenarios. In the first, a group of bandits enters a village and threatens to kill any inhabitant who does not hand over their money; once the inhabitants comply, the bandits leave in search of another village to rob. In the second, the situation is the same, except that the bandits return occasionally, though irregularly. In the third, the bandits set a permanent establishment in the village that is meant to suppress any hint of unrest. With the money they accumulate, the bandits are able to progressively enlarge their contingent. Over time, they no longer need to threaten

anyone, since each inhabitant simply prepares the money in advance and a bandit comes to collect it. In all three scenarios, the two conditions identified above appear to be satisfied. The bandits' actions alter the inhabitants' possible strategies and program for a particular outcome in a counterfactually robust way: under normal circumstances, i.e., assuming the inhabitants value their lives, handing over the money is their only rational strategy, regardless of the specific reasons behind individuals' preferences or the details of their decision-making (e.g., whether they want to stay alive for the sake of their family or in pursuit of great achievements). Intuitively, however, it is only the third scenario that counts as a structural change to the village's political or economic structure. What explains this difference?

To find an answer, we can look at a third distinctive feature of social structures. Several authors have argued that social structures exhibit a significant degree of stability. They are more stable, in particular, than individual coercive threats. This point is explicit in Iris Marion Young, who notes that social structures "appear as objective, given, and constraining. [They] do not constrain in the form of the direct coercion of some individuals over others; they constrain more indirectly and cumulatively as blocking possibilities" (2011: 55). Lauren Ross similarly argues that "structural factors are often assumed to be fixed, unchanging, and stable. This is because structure changes less often (or on longer time scales) and it can be more difficult to change this factor relative to others" (175).

These remarks point in the right direction, but they are unsatisfactory as stated. Against Young, one may note that social structures *can* be coercive in a relatively direct way, as in our third scenario. Against Ross, one may note that fixity and invariance are not sufficient to distinguish structural factors from coercive threats: certain intuitively non-structural factors may be invariant for purely contingent reasons, as in our second scenario, where the bandits' threat is recurrent yet does not count as a structural factor. The stability of social structures, by contrast, is not a mere accident. Rather, social structures have features that ensure their own stability. To understand the stability of social structures, we therefore need to clarify what these features are.

Our view is that social structures are stable insofar as, and because, they generate *feedback loops* that sustain the outcome the structure programs for, and by so doing, create conditions under which preserving the structure is more advantageous for the relevant stakeholders than challenging it, thereby supporting and reinforcing the structure's continued existence.

Let us unpack the main elements of this characterization. Feedback loops can be of two kinds: self-reinforcing (or positive) feedback loops and balancing (or negative) feedback loops. A self-reinforcing loop is a process in which a factor enables or produces effects that elicit responses that amplify the initial effects. As an example, consider how 'likes' on social media platforms contributes to virality by increasing visibility and prompting further engagement, which in turn reinforces the post's original exposure. A balancing loop, by contrast, is a process in which deviations from a targeted state trigger responses that return the system to that state, as when a thermostat activates in response to temperature diverging from the target in order to restore it.

Social structures involve both kinds of feedback loops. These feedback loops sustain the outcome the social structure programs for. However, they do so in different ways. Self-reinforcing feedback loops *strengthen* the outcome by ensuring that both the type and the direction of the effects the social structure enables or produces remain consistent. By contrast, balancing feedback loops

stabilize the outcome by counteracting deviations and maintaining the system within a targeted range.

When such feedback loops are in place, certain strategies become the predictable, default, and rational course of action for significant groups of agents occupying a range of social positions. As a result, a class of relevant stakeholders emerges for whom the advantages of having the structure are superior to the costs, and who thereby become depository of resources, energies, and dispositions that ensure or make it more likely that, were the existence of the current structure to be challenged, they would resist the change. Of course, not all agents benefit from the persistence of social structures. Yet the feedback loops these structures generate often make the role of ‘agent of change’ prohibitively costly even for those that are disadvantaged by them. Under such conditions, preserving the structure is typically preferable to attempting to transform it. In this way, the feedback loops generated by the structure produce the internal conditions that secure its own stability.

To put things differently by reference to structural factors, self-reinforcing and balancing feedback loops typically lead to something like lock-in effects, where each iteration reinforces the role and importance of the factors that produce such feedback loops within the system, making it increasingly difficult or costly to alter them. This reinforcement can occur in two, often complementary, ways. Each iteration may produce increased benefits for relevant stakeholders, which favor continued reliance on the factor. Alternatively, each iteration may increase the disadvantage for the individuals involved of adopting different strategies or modifying the factor in question, discouraging change. In both cases, the factor becomes ‘entrenched’ within the system. By so doing, these mechanisms ensure the stability of the structural factor itself.

These considerations lead us to a third necessary condition for structural change: for a change to count as structural, it must generate feedback loops that sustain the outcome the change programs for and, by so doing, create conditions that support and reinforce the factor through which the change occurs. Since any change can be redescribed in terms of the factor introduced and the outcome it produces, we can restate this condition more compactly: for a change to count as structural, it must generate feedback loops that make the change self-sustaining.

We can illustrate the third condition by examining how it operates both in our earlier example and in the paradigmatic cases of structural change discussed above. The bandit’s third scenario combines both kinds of feedback loops. By establishing a permanent contingent, the bandits create a balancing feedback loop that counteracts deviations and stabilizes the outcome their threat programs for: the regular handing over of money. But this is not all: their contingent also generates a self-reinforcing feedback loop. The money they extract lets them enlarge their contingent, which in turn lets them extract still more money and tighten their grip on the village; with each iteration, their hold over its political and economic life becomes further entrenched. Through this repeated cycle, the bandits’ threat, though no longer explicitly issued, becomes self-sustaining. Nothing of this sort happens in the first two scenarios – no feedback loop entrenches the bandits’ presence there – which explains our intuition that neither scenario counts as a structural change.

Turning now to the paradigmatic cases: as we have seen, when the Supreme Court allowed corporations and wealthy donors to spend unlimited sums on elections, it became rational for candidates across the political spectrum to actively seek corporate support. In turn, the growing

role of such donors in determining electoral outcomes enhanced their bargaining power, making it especially advantageous for them to invest in elections. As contributions increased, politicians became more reliant on corporate backing, which further strengthened donors' position and incentives to invest in elections, which made politicians even more reliant on them, and so on. This self-reinforcing cycle has not only made the effects of the ruling persistent over time; it has also progressively entrenched and solidified the role of the new campaign finance laws in shaping American politics. As more politicians are elected through dark money and become beholden to corporations and wealthy donors, it becomes increasingly difficult to reform a system of campaign finance laws that favors the political dominance of the wealthiest.

A similar dynamic is at work in the case of the European Single Market. By making increased cross-border trade a rational strategy for firms across member states, the Single Market generated cycles that boosted, in aggregate, economic growth and consumer benefits. This gradually increased economic interdependence among individuals and firms within and across European countries, reinforcing both the practical importance of the single market and the legal framework that sustains it. The more deeply integrated European economies became, the higher the costs of dismantling the arrangement, and the stronger the incentives to preserve and extend it.¹⁸

While *Citizens United* and *European Single Market* are examples of changes that generated self-reinforcing feedback loops, the Federal Reserve was explicitly designed to serve as a balancing mechanism, preventing runaway effects such as deepening recessions or systemic banking crises. In general, the output gap – the difference between actual real GDP and potential output – is the variable that triggers the Fed's intervention in interest rate control. When the output gap is negative (i.e., actual GDP is below potential output), the Fed implements expansionary monetary policy, lowering interest rates to stimulate economic activity. Conversely, when the output gap is positive (i.e., the economy is overheating), the Fed adopts contractionary monetary policy, raising interest rates to curb inflation and excessive growth. Through these interventions, the Fed aims to keep the economy within a stable and sustainable range. The creation of the Federal Reserve has changed the strategic landscape of the relevant actors in such a way that is almost unthinkable to imagine a scenario in which the Federal Reserve could be abolished – although, as the current state of affairs demonstrates, this does not prevent attempts to change a fundamental aspect of the institutions, namely, its independence. Consistent with our account, however, these attempts have been met with significant resistance from a wide range of relevant stakeholders.

We want to conclude this subsection by clarifying the distinction between the second and the third condition and by addressing two potential objections to the third condition. As explained before, the second condition concerns the robustness of the outcome a change produces, whereas the third concerns the stability of the factor through which the change occurs. Though distinct, the two conditions are similar in some respects. To see why, consider the following redescriptions. According to the second condition, a structural change requires introducing a factor that makes certain strategies rational for the individuals involved, regardless of who these individuals are and what motivates them; by so doing, the factor ensures, or makes it more likely, that a particular outcome will occur, in one way or another. In this way, the factor that is introduced 'programs' for a particular outcome. By contrast, according to the third condition, a structural change requires

¹⁸ A clear sign of this self-reinforcing dynamic is the significant legal and economic challenges faced by Brexit proponents in their attempts to disentangle the UK from the European Union.

introducing a factor that generates self-reinforcing or balancing feedback loops, as a result of which it becomes rational for the individuals involved to pursue strategies that preserve or maintain the factor – and this, regardless of who these individuals are and, indeed, even when they do not benefit from the outcome it produces. An attentive reader will notice that, in so doing, the factor effectively ‘programs’ for its own sustainment and persistence. Thus, both the second and the third conditions hold that for a change to count as structural, the factor through which the change occurs must have programming features. What differs is the object of that programming. The second condition requires the relevant factor to program for a particular outcome; the third requires it to program for its own persistence.

This quite naturally leads us to two potential objections. The first is that requiring a structural change to be stable in the way described seems to make the idea of a self-destructive structure conceptually impossible. This would be an unwelcome implication of our account, given that such ideas play a central role in various philosophical traditions. For instance, Marxist critiques often argue that capitalism is a system that generates the conditions for its own downfall. We believe, however, that our account is fully compatible with the notion of a self-destructive structure. A structure typically consists of multiple structural factors, each with its own feedback dynamics, and these dynamics may be in tension with one another. Internal contradictions of this kind can generate pressures that ultimately lead to the structure’s collapse. That said, if it is a genuine structure, its collapse will meet resistance rather than occur outright, because it is part of the constitutive features of a structure to involve feedback loops that make it resilient to change. This is precisely what we observe in the case of capitalist systems. More broadly, our emphasis on stability should not be taken to imply that structural changes produce static or unchanging structures. Social structures evolve. A structural change may set in motion dynamics that transform social relations and practices over time, generating new – possibly structural – problems that call for further structural change.

The second objection is that some intuitively structural changes appear to lack the required stability. Consider a new green technology, e.g., a cheaper source of renewable energy, which alters the strategic landscape by making its adoption the more advantageous consumption strategy, thereby programming for a specific outcome and beginning to generate self-reinforcing feedback loops. Suppose this technology is swiftly banned by an administration beholden to the coal industry. The technology appears to satisfy our three conditions, yet the change it initiates is deeply unstable. In response, we want to emphasize that, on our account, for the introduction of a factor to be a structural change, it is not necessary that the factor itself *actually* persist over time. Rather, the factor must *program* for its own persistence, that is, it must have features that ensure, *or make it more likely*, that it will persist across a range of variations. This is, however, compatible with the possibility that external circumstances might prevent its programming features from manifesting. Consider a particular economic organization that generates feedback loops of the relevant sort. Such an organization may fail to persist if, say, a meteorite destroys the country’s main economic infrastructure. Because the organization has the relevant programming features, it counts as stable in our sense, even though it does not actually persist because the meteorite prevents its program from running, so to speak. In other words, the question the third condition asks is not whether a factor generates feedback loops that make it capable of resisting *any* external shock; it is whether the factor generates feedback loops that would make it self-sustaining, rather than having to depend entirely on continued external support and enforcement for its existence. The green technology described above has the capacity to generate the kind of dynamics that would make it self-

sustaining – indeed, its feedback loops had already begun to operate before the ban cut them short. Its introduction therefore still counts as a structural change, even though external circumstances prevented that change from persisting.

3.4. *Applying the Account*

Let us sum up the results of our analysis. We have argued that for a change to qualify as structural, it must satisfy the following conditions.

- (1) *The strategy set condition*: It must alter the set of possible strategies open to the individuals involved.
- (2) *The robustness condition*: It must program for a new outcome or for an outcome that was previously unprogrammed.
- (3) *The stability condition*: It must generate feedback loops that make the change self-sustaining.

One way to test our account of structural change is to apply it to uncertain cases and examine whether it delivers verdicts that are both theoretically compelling and that match our considered judgments. *Québec Bill 83* serves this purpose well. Recall that the Bill is intended to transform the structure of Québec's health system, which is currently shaped by incentives that encourage Québec-trained doctors to leave the public sector. In our terms, the present structure programs for an outcome characterised by a progressive depletion of doctors in the public system, and does so in a self-reinforcing way: the fewer doctors remain, the more difficult the working conditions become, and the stronger the incentive to leave. Changing this structure requires introducing a factor that programs for a different outcome, namely, a stable supply of doctors within the public system, and that is self-sustaining. We think that Bill 83 fails to achieve this.

Consider first what Bill 83 actually achieves. While the Bill does not technically alter the *actions* available to Québec-trained doctors (they may still choose to work in Québec's public system, move to another province, or enter the private sector), it clearly modifies their *options* by imposing severe financial penalties on any choice other than remaining in the public system during the first five years of a doctor's career. These penalties are so heavy that they render the alternatives non-genuine options, since the outcomes associated with the technically still-available actions now involve prohibitive costs. By so doing, Bill 83 also reshapes the set of possible career *strategies* open to new doctors, effectively ensuring that most spend their first five years in Québec's public health system. In this respect, the Bill satisfies the first two conditions: it alters the doctors' strategy set and programs for a particular outcome.

It fails, however, on the third condition. To begin with, the feedback loops generated by the Bill are limited. Once the five-year obligation expires, the system contains no internal mechanism for retaining doctors. As long as the underlying incentives to work elsewhere remain intact (e.g., better conditions, higher pay, greater freedom to choose one's location), leaving the public system remains the long-term rational strategy, and doctors are likely to do so as soon as they are free, much as a captive flees once a threat is lifted. The system remains vulnerable to attrition, particularly when compounded by predictable retirements, since nothing within it counteracts the drain of doctors once the mandatory period ends. In sum, Bill 83 does not modify the internal

dynamics of the system in a way that would ensure a stable supply of doctors in public health over time; it functions as a temporary coercive constraint rather than as a balancing mechanism.

More decisively, Bill 83 does not create the conditions that would make it self-sustaining. Nothing in the Bill ensures, or makes it more likely, that it will become entrenched in the health system. On the contrary, the Bill appears to depend on nothing more than the government's willingness to maintain it in place and fund its enforcement. If anything, we should expect that those with the resources and influence to seek its abolition, i.e., the powerful doctor lobbies chief among them, will do just that. Nor can the government count on the support of the general population: as we have seen, the Bill fails to prevent attrition in the public health system, so the benefits capable of mobilizing the population to defend the Bill will be correspondingly limited. For all these reasons, the Bill does not count as a structural change on our account – a verdict we regard as deeply plausible.

4. Conclusion

At times, the social structure is so fundamentally transformed by the introduction of certain factors, e.g., laws, policies, organizations, technologies, that the change is best described as a structural change. In some cases, these changes are precisely what is required to address structural problems.

In this paper, we have developed an account that helps identify instances of structural changes, distinguish them from non-structural changes, and explain why they produce the effects they do. Our account also offers practical guidance for bringing about structural changes: to enact them, one must identify a factor, or combination of factors, whose introduction alters the possible strategies open to the relevant agents and makes some of these strategies rational for them in a way that is both robust and stable. Our account therefore satisfies the desiderata we identify at the outset: it is descriptively and explanatorily adequate as well as action-guiding.

To offer a complete theory of how to address structural problems, further work is required. However, our account already achieves two things. First, it provides a tool for evaluating the merits of policies and innovations as candidate structural solutions to structural problems. Insofar as structural solutions require structural changes, our account suggests that common public policies that aim to address structural problems merely by adjusting incentives, such as income tax adjustments or fiscal credits, are not solutions of the right kind, because they do not themselves constitute structural changes. They modify the parameters of elements already present in the structure, without altering the possible strategies open to the relevant agents. Likewise, policies that depend on specific motivations, such as training programs meant to improve online behaviour, or that rely entirely on the state's external support for their continued existence, such as price controls to reduce inflation, fail to count as structural solutions, since they lack, respectively, the relevant programming features and the relevant feedback loops that would make them structural changes.

Second, and more importantly, our account offers the basis for the development of a broader theory of how to address structural problems through structural solutions. The very conditions we have identified for a change to count as structural can indeed be used to understand what makes a problem, and a solution to it, structural. At the most basic level, a structural problem is a problem

that is caused and sustained by a particular social structure. But since a structure operates as it does because of the structural factors that constitute it, our account suggests a more precise characterization: a structural problem is a problem that is caused and sustained by a structural factor, or combination of structural factors – that is, by factors that play a central role in determining the possible strategies open to the relevant agents and that, partly by doing so, program for a problematic outcome and generate the kind of feedback loops that sustain both that outcome and the factor itself. A solution to a structural problem so conceived must address the structural factors responsible for the problem, and it must do so in a way that brings about genuine structural change. Developing this account in detail is both important and urgent. It is the task we have taken on in a companion paper.¹⁹

Acknowledgments: We are very grateful to Sean Aas, Arash Abizadeh, Manula Adhihetty, David Axelsen, Paul Bou-Habib, Pia Campeggiani, Dulyaphab Chaturongkul, James Christensen, Andrée-Anne Cormier, Samuel Delorme, Sally Haslanger, Richard Healey, Colin Macleod, Sebastiano Moruzzi, Angie Pepper, João Rosas, Gianni Sarra, Jordan Scott, Christine Tappolet, Isabella Trifan, Andy Yu, Raphael van Riel, Giorgio Volpe, the participants to conferences at the University of Bologna, the *Essex Political Theory Work-in-Progress Workshop*, the *Canadian Philosophical Association 2026 Congress* at Dalhousie University, the *16th Braga Meeting on Ethics and Political Philosophy*, the *2026 Conference of the Association for Social and Political Philosophy* at St Andrews University, and the *2nd International Meeting of the PPE Society* at King’s College London, for their very helpful comments on previous versions and presentations of this paper.

References

Authors 1&2.

Collins, S. (2025). Rage Against the Machine. *Ethics*, 135 (4): 704–717.

Estlund, D. (2024). What’s Unjust about Structural Injustice? *Ethics*, 134 (3): 333–359.

Estlund, D. (2026). Injustice and Structural Wrong: Reply to Sankaran and Monaghan, and Collins. *Ethics*, 136 (3): 654–669.

Haslanger, S. (2016). What is a (social) structural explanation? *Philosophical Studies*, 173: 113–130.

Haslanger, S. (Manuscript). How to Change a Social Structure.

Hindriks, F. (2025). *The Structure of the Open Society: Social Ontology Meets Collective Ethics*. Oxford: Oxford University Press.

Jackson, F. & Pettit, P. (1992). Structural explanation in social theory. In D. Charles and K. Lennon (Eds.), *Reduction, explanation and realism*, pp. 97–131. Oxford University Press.

McKeown, M. (2021). Structural injustice. *Philosophy Compass*, 16 (7): e12757.

McKeown, M. (2024). *With Power Comes Responsibility: The Politics of Structural Injustice*. Bloomsbury Academic.

Ross, L. N. (2024). What is social structural explanation? A causal account. *Noûs*, 58 (1):163–179.

Sangiovanni, A. (2018). Structural Injustice and Individual Responsibility. *Journal of Social Philosophy*, 49 (3): 461–483.

¹⁹ See [REFERENCE DELETED].

- Sankaran, K. (2021). “Structural Injustice” as an analytical tool. *Philosophy Compass*, 16 (10): e12780.
- Sankaran, K. & Monaghan, J. (2025). Putting the ‘Structural’ Back in ‘Structural Injustice’. *Ethics*, 135 (3): 545–559.
- 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.